arXiv:1512.06798v1 [math.PR] 21 Dec 2015

LIMITS OF DISCRETE DISTRIBUTIONS AND GIBBS MEASURES ON RANDOM GRAPHS

AMIN COJA-OGHLAN*, WILL PERKINS, KATHRIN SKUBCH

ABSTRACT. Building upon the theory of graph limits and the Aldous-Hoover representation and inspired by Panchenko's work on asymptotic Gibbs measures [Annals of Probability 2013], we construct continuous embeddings of discrete probability distributions. We show that the theory of graph limits induces a meaningful notion of convergence and derive a corresponding version of the Szemerédi regularity lemma. Moreover, complementing recent work [Bapst et. al. 2015], we apply these results to Gibbs measures induced by sparse random factor graphs and verify the "replica symmetric solution" predicted in the physics literature under the assumption of non-reconstruction.

Mathematics Subject Classification: 05C80, 82B44

1. INTRODUCTION

The systematic study of limits of discrete structures such as graphs or hypergraphs emerged about a decade ago [13, 29]. It has since become a prominent and fruitful area of research, with numerous applications in combinatorics and beyond. The basic idea is to embed discrete objects into a "continuous" space so that tools from analysis, topology and measure theory can be brought to bear. Conversely, the connection extends genuinely combinatorial ideas such as the Szemerédi regularity lemma to the continuous world.

In this paper we study "analytic embeddings" of probability distributions on discrete cubes, i.e., sets of the form Ω^V for some fixed finite set Ω and a large finite set V. Arguably, probability measures on discrete cubes are among the most important and most basic objects in combinatorics, computer science and mathematical physics. For instance, they occur as the Gibbs measures of finite spin systems such as the Ising model [31]. In this case $\Omega = \{\pm 1\}$ and V is a finite set of lattice points, whose size ultimately goes to infinity in the "thermodynamic limit". Similarly, much of the theory of Markov Chain Monte Carlo deals with the correlations that, e.g., the uniform distribution on the set of k-colorings of some graph $G = (V, E)$ induces [28]. Thus, Ω would be the set $[k] = \{1, \ldots, k\}$ and V is the vertex set of a (large) graph. We are going to construct limiting objects of such measures, state a corresponding regularity lemma and illustrate applications to random graphs.

Perhaps the most prominent related construction is the Aldous–Hoover representation of exchangeable arrays [5, 21]. Its connection to graph limits was observed by Diaconis and Janson [19]. Furthermore, Panchenko [33] used the Aldous-Hoover representation to introduce the notion of "asymptotic Gibbs measures", the protagonists of his work on "mean-field models" such as the Sherrington–Kirkpatrick model. Indeed, Panchenko has a proof of the Aldous-Hoover result from graph limits [34, Appendix A]. Our embedding of discrete measures into continuous space can be seen as a generalization of this approach to arbitrary measures on discrete cubes.

The main contributions of the present work are as follows. First, in Section 2 we construct a natural embedding of general measures on discrete cubes into a "continuous" metric space. We highlight connections to the Aldous–Hoover representation and the theory of graph limits. Moreover, we state a "regularity lemma" for such general measures that is a bit stronger than the regularity lemma of Bapst and Coja-Oghlan [8]. In a sense, the main point of Section 2 is to study discrete probability measures using some of the main ideas from the theory of graph limits. Second, building upon ideas of Panchenko [33], in Section 3 we apply the concepts of Section 2 to Gibbs measures induced by sparse random graphs. In particular, we verify the "replica symmetric solution" predicted in the physics literature [26] under the assumption that the sequence Gibbs measures converges in probability as the size n of the random graph tends to infinity. Additionally, we will see that the sequence of Gibbs measures induces a "geometric" limiting object on the (infinite) Galton-Watson tree that describes the local structure of the sparse random graph. Further, we show that a spatial mixing property called non-reconstruction is a sufficient condition for "replica symmetry".

*The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC Grant Agreement n. 278857–PTCC.

Many of the proofs build upon known methods, although the proofs in Section 3 require quite a bit of technical work. Specific references will be given as we proceed, but let us point to Lovász' comprehensive treatment [29] of graph limits and to Janson's work [22] that provides some of the measure-theoretic foundations.

Notation and preliminaries. For a measurable space $(\Omega,\mathscr{F})$ we denote the set of probability measures by $\mathscr{P}(\Omega,\mathscr{F})$ or briefly by $\mathscr{P}(\Omega)$. Moreover, for $x\in\Omega$ we denote by $\delta_x\in\mathscr{P}(\Omega)$ the Dirac measure on x. Further, we write $\lambda(\cdot)$ for the Lebesgue measure on $\mathbb{R}$. If Ω is a finite set, then the σ-algebra is always understood to be its power set. Given $\mu\in\mathscr{P}(\Omega^n)$, we write $\boldsymbol{\sigma},\boldsymbol{\sigma}',\ldots$ for independent samples from μ. Additionally, if $X:\Omega^n\to\mathbb{R}$ is a function, then $\langle X(\boldsymbol{\sigma})\rangle_\mu=\sum_{\sigma\in\Omega^n}X(\sigma)\mu(\sigma)$ signifies the mean of X.

We often use the following notation to define a probability measure p on a finite set Ω. If $f:\Omega\to[0,1]$ is a function that is not identically 0, then $p(\omega)\propto f(\omega)$ is a shorthand for

$$p(\omega)=f(\omega)\Big/\sum_{\omega'\in\Omega}f(\omega').$$

We shall frequently work with the spaces $L_1(U,\mathbb{R}^l)$ for a measurable $U\subset\mathbb{R}^k$ (with k,l natural numbers). Recall that this is the space of measurable functions f such that $\int|f|<\infty$, up to equality almost everywhere. We tacitly identify the elements of $L_1(U,\mathbb{R}^l)$ with specific (fixed) representatives, i.e., measurable functions $U\to\mathbb{R}^l$, so that we can write $f(x)$ for $x\in U$. Moreover, we remember that a *Polish space* is a complete metric space that has a countable dense subset. Examples include the spaces $\mathbb{R}^k$ and $L_1(U,\mathbb{R}^l)$ for any k,l.

If $\mu\in\mathscr{P}(\Omega)$ and $\nu\in\mathscr{P}(\Omega')$, then we write $\Gamma(\mu,\nu)$ for the set of all couplings of μ,ν. Thus, $\Gamma(\mu,\nu)$ is the set of all a probability measures $\gamma\in\mathscr{P}(\Omega\times\Omega')$ such that $(x,y)\in\Omega\times\Omega'\mapsto x$ maps γ to μ and $(x,y)\in\Omega\times\Omega'\mapsto y$ maps γ to ν. Moreover, $S_{[0,1)}$ is the set of all measurable $f:[0,1)\to[0,1)$ such that $f(\lambda)=\lambda$. In addition, for a random variable $X:(\Omega,\mu)\to\Omega'$ we write $\mathscr{L}(X)=X(\mu)$ for the distribution of X, i.e., the probability measure on Ω' defined by $A'\mapsto\mu(X^{-1}(A'))$ for measurable $A'\subset\Omega'$. Finally, if μ is a probability measure on a product space Ω^V and $U\subset V$, then $\mu_{\downarrow U}$ signifies the marginal distribution of μ on the coordinates U. If $U=\{u\}$ is a singleton, then we write $\mu_{\downarrow u}=\mu_{\downarrow\{u\}}$.

We recall that for probability measures μ,ν defined on a metric space E with metric $D(\cdot,\cdot)$ the L_1*-Wasserstein metric* is defined as

$$d_1(\mu,\nu)=\inf\biggl\{\int_{\Omega\times\Omega}D(x,y)\,d\gamma(x,y):\gamma\in\Gamma(\mu,\nu)\biggr\}.$$

For probability measures on a compact Polish space this metric induces the topology of weak convergence.

We are going to use the following well-known property of the Poisson distribution, and call it the "Chen-Stein" property after [15]. If X has distribution $\mathrm{Po}(d)$ for some $d>0$ and if $f(X)$ is a function such that $\mathrm{E}[X|f(X)|]<\infty$, then

$$\mathrm{E}[Xf(X)]=d\mathrm{E}[f(X+1)] \tag{1.1}$$

2. Probability measures on cubes

2.1. **The cut metric.** Fix a finite set Ω and let $n\geq1$ be an integer. We would like to identify a measure $\mu\in\mathscr{P}(\Omega^n)$ with a "continuous object". To this end, we represent a point $\sigma=(\sigma_1,\ldots,\sigma_n)\in\Omega^n$ by the function

$$\hat\sigma:[0,1)\to\mathscr{P}(\Omega),\qquad x\mapsto\sum_{i=1}^n\delta_{\sigma_i}\mathbf{1}\{x\in[(i-1)/n,i/n)\}.$$

Thus, $\hat\sigma$ is a step function whose value on the interval $[(i-1)/n,i/n)$ equals the Dirac measure $\delta_{\sigma_i}\in\mathscr{P}(\Omega)$. Because Ω is finite, $\mathscr{P}(\Omega)$ is just a simplex of dimension $|\Omega|-1$ in $\mathbb{R}^\Omega$. The induced Borel σ-algebra turns $\mathscr{P}(\Omega)$ into a Polish space. Moreover, $\mu\in\mathscr{P}(\Omega^n)$ corresponds to the probability measure $\hat\mu$ on $L_1([0,1],\mathbb{R}^\Omega)$ defined by

$$\hat\mu=\langle\delta_{\hat{\boldsymbol{\sigma}}}\rangle_\mu=\sum_{\sigma\in\Omega^n}\mu(\sigma)\delta_{\hat\sigma}. \tag{2.1}$$

Indeed, let $\Sigma_\Omega=L_1([0,1),\mathscr{P}(\Omega))\subset L_1([0,1),\mathbb{R}^\Omega)$ be the space of all measurable functions $f:[0,1)\to\mathscr{P}(\Omega)$ with values in $\mathscr{P}(\Omega)$ up to equality almost surely. Then Σ_Ω is Polish and $\hat\mu\in\mathscr{P}(\Sigma_\Omega)$. Clearly, the map $\mu\mapsto\hat\mu$ is one-to-one. Further, by extension of the discrete notation, we denote the mean of a measurable X on Σ_Ω with respect to $\mu\in\mathscr{P}(\Sigma_\Omega)$ by

$$\langle X(\boldsymbol{\sigma})\rangle_\mu=\int_{\Sigma_\Omega}X(\sigma)\,d\mu(\sigma).$$

Then the continuous analogue of (2.1) reads $\mu=\langle\delta_{\boldsymbol{\sigma}}\rangle_\mu$.

Following the theory of graph limits [29], we define the *strong cut metric* as

$$D_\square(\mu,\nu) = \inf_{\gamma\in\Gamma(\mu,\nu)} \sup_{B\subset\Sigma_\Omega^2, U\subset[0,1)} \left\| \int_B \int_U \sigma_x - \tau_x \, dx \, d\gamma(\sigma,\tau) \right\|_1 \qquad (\mu,\nu\in\mathscr{P}(\Sigma_\Omega)), \tag{2.2}$$

where, of course, B, U are understood to be measurable. Additionally,

$$\Delta_\square(\mu,\nu) = \inf_{\gamma\in\Gamma(\mu,\nu), s\in S_{[0,1)}} \sup_{B\subset\Sigma_\Omega^2, U\subset[0,1)} \left\| \int_B \int_U \sigma_x - \tau_{s(x)} \, dx \, d\gamma(\sigma,\tau) \right\|_1 \qquad (\mu,\nu\in\mathscr{P}(\Sigma_\Omega)), \tag{2.3}$$

is the *weak cut metric*. The general results [22] imply

Fact 2.1. *$D_\square(\cdot,\cdot)$ is a metric and $\Delta_\square(\cdot,\cdot)$ is a pseudo-metric on $\mathscr{P}(\Sigma_\Omega)$. Moreover, the infima in (2.2)–(2.3) are attained.*

Let us write M_Ω for the space $\mathscr{P}(\Sigma_\Omega)$ endowed with the metric $D_\square(\cdot,\cdot)$. Moreover, call $\mu,\nu\in M_\Omega$ *equivalent* if $\Delta_\square(\mu,\nu)=0$ and write $\bar\mu$ for the equivalence class of $\mu\in M_\Omega$. Then $\Delta_\square(\cdot,\cdot)$ induces a metric on the space $\mathscr{M}_\Omega$ of equivalence classes. We shall see momentarily that $\mathscr{M}_\Omega$ (essentially) coincides with the usual graphon space. But let us first look at two examples.

Example 2.2. *Let $\Omega=\{0,1\}$, $p=\mathrm{Be}(1/2)$ and let $\mu_n = p^{\otimes n}\in\mathscr{P}(\Omega^n)$ be the uniform distribution on the Hamming cube. Letting $\sigma:[0,1)\to\mathscr{P}(\Omega)$, $x\mapsto p$ be constant, we expect that μ_n converges to $\nu=\delta_\sigma$ as $n\to\infty$. Indeed, because ν is a Dirac measure there is just one coupling γ of μ,ν. Hence, the set B in (2.2) really boils down to a set $B_0\subset\Omega^n$ of configurations and the set $U\subset[0,1]$ to a weight function $u:[n]\to[0,1/n]$ such that*

$$\int_B\int_U p-\sigma(x)\,dx\,d\gamma(\sigma,\tau) = p - \sum_{\sigma\in B_0}\sum_{x=1}^n u(x)\sigma(x). \tag{2.4}$$

Since for $\boldsymbol\sigma$ chosen from μ the sum $\sum_{x=1}^n u(x)\boldsymbol\sigma(x)$ comprises of independent summands, Azuma's inequality shows that for there is a constant $c>0$ such that for any $t>0$ $\mathrm{P}\big[\big\|p-\sum_{x=1}^n u(x)\boldsymbol\sigma(x)\big\|_1 > t/\sqrt{n}\big] < 2\exp(-ct^2)$. Therefore, the norm of (2.4) is $O(n^{-1/2})$ for all B_0, U.

Example 2.3. *Let $\Omega=\{0,1\}$, $p=\mathrm{Be}(1/3)$, $q=\mathrm{Be}(2/3)$ and $\mu=\frac12(p^{\otimes n/2}\otimes q^{\otimes n/2}+q^{\otimes n/2}\otimes p^{\otimes n/2})$ for an even $n>1$. Let $\sigma:[0,1)\to\mathscr{P}(\Omega)$, $x\mapsto p\mathbf{1}\{x<1/2\}+q\mathbf{1}\{x\ge1/2\}$ and $\tau:[0,1)\to\mathscr{P}(\Omega)$, $x\mapsto\sigma(1-x)$. Let $\nu=\frac12(\delta_\sigma+\delta_\tau)$. Then*

$$\Delta_\square(\mu,\nu) = D_\square(\mu,\nu) = O(n^{-1/2}). \tag{2.5}$$

Indeed, to construct a coupling γ of μ,ν let X,Y,Y' be three independent random variable such that $X=\mathrm{Be}(1/2)$, $Y\in\{0,1\}^n$ has distribution $p^{\otimes n/2}\otimes q^{\otimes n/2}$ and $Y'\in\{0,1\}^n$ has distribution $p^{\otimes n/2}\otimes q^{\otimes n/2}$. Further, let $G=(Y,\delta_\sigma)$ if $X=0$ and $G=(Y',\delta_\tau)$ otherwise and let γ be the law of G. A similar application of Azuma's inequality as in the previous example yields (2.5).

2.2. **Alternative descriptions.** We recall that the (bipartite, decorated version of the) cut metric on the space W_Ω of measurable maps $[0,1)^2\to\mathscr{P}(\Omega)$ can be defined as

$$\delta_\square(f,g) = \inf_{s,t\in S_{[0,1)}} \sup_{U,V\subset[0,1)} \left\| \int_{U\times V} f(x,y) - g(s(x),t(y))\,dx\,dy \right\|_1 \qquad \text{(cf. [22, 29, 30])}.$$

Let $\mathscr{W}_\Omega$ be the space obtained from W_Ω by identifying $f,g\in W_\Omega$ such that $\delta_\square(f,g)=0$. Applying [22, Theorem 7.1] to our setting, we obtain

Proposition 2.4. *There is a homeomorphism $\mathscr{M}_\Omega\to\mathscr{W}_\Omega$.*

Proof. We recall that for any $\mu\in\mathscr{P}(\Sigma_\Omega)$ there exists a measurable $\varphi:[0,1)\to\Sigma_\Omega$ such that $\mu=\varphi(\lambda)$, i.e., $\mu(A)=\lambda(\varphi^{-1}(A))$ for all measurable $A\subset\Sigma_\Omega$. Hence, recalling that $\varphi(x)\in L_1([0,1),\mathscr{P}(\Omega))$, μ yields a graphon $w_\mu:[0,1]^2\to\mathscr{P}(\Omega)$, $(x,y)\mapsto(\varphi(x))(y)$. Due to [22, Theorem 7.1] the map $\bar\mu\in\mathscr{M}_\Omega\mapsto w_\mu\in\mathscr{W}_\Omega$ is a homeomorphism. □

Corollary 2.5. *$\mathscr{M}_\Omega$ is a compact Polish space.*

Proof. This follows from Proposition 2.4 and the fact that $\mathscr{W}_\Omega$ has these properties [29, Theorem 9.23]. □

Diaconis and Janson [19] pointed out that the connection between $\mathscr{W}_\Omega$ and the Aldous-Hoover representation of "exchangeable arrays" (see also Panchenko [34, Appendix A]). To apply this observation to $\mathscr{M}_\Omega$, recall that $\Omega^{\mathbb{N}\times\mathbb{N}}$ is compact (by Tychonoff's theorem) and that a sequence $(A(n))_n$ of $\Omega^{\mathbb{N}\times\mathbb{N}}$-valued random variables converges to A in distribution iff

$$\lim_{n\to\infty} \mathrm{P}\left[\forall i,j\le k : A_{ij}(n) = a_{ij}\right] = \mathrm{P}\left[\forall i,j\le k : A_{ij} = a_{ij}\right] \qquad \text{for all } k,\ a_{ij}\in\Omega.$$

Now, for $\bar\mu\in\mathscr{M}_\Omega$ define a random array $\boldsymbol{A}(\bar\mu) = (\boldsymbol{A}_{ij}(\bar\mu))\in\Omega^{\mathbb{N}\times\mathbb{N}}$ as follows. Let $(\boldsymbol{\sigma}_i)_{i\in\mathbb{N}}$ be a sequence of independent samples from the distribution μ, independent of the sequence $(\boldsymbol{x}_i)_{i\in\mathbb{N}}$ of independent uniform samples from $[0,1)$. Finally, independently for all i,j choose $\boldsymbol{A}_{ij}(\bar\mu)\in\Omega$ from the distribution $\boldsymbol{\sigma}_i(\boldsymbol{x}_j)\in\mathscr{P}(\Omega)$. Then in our context the correspondence from [19, Theorem 8.4] reads

Corollary 2.6. *The sequence $(\bar\mu_n)_n$ converges to $\bar\mu\in\mathscr{M}_\Omega$ iff $\boldsymbol{A}(\bar\mu_n)$ converges to $\boldsymbol{A}(\bar\mu)$ in distribution.*

While Corollary 2.6 characterizes convergence in $\Delta_\square(\cdot,\cdot)$, the following statement applies to the strong metric $D_\square(\cdot,\cdot)$. For $\sigma\in\Sigma_\Omega$ and $x_1,\ldots,x_k\in[0,1)$ define $\sigma_{\downarrow x_1,\ldots,x_k} = \sigma(x_1)\otimes\cdots\otimes\sigma(x_k)\in\mathscr{P}(\Omega^k)$. Moreover, for $\mu\in M_\Omega$ let

$$\mu_{\downarrow x_1,\ldots,x_k} = \int_{\Sigma_\Omega} \sigma_{\downarrow x_1,\ldots,x_k}\,d\mu(\sigma).$$

If $\mu\in\mathscr{P}(\Omega^n)$ is a discrete measure, then $\hat\mu_{\downarrow x_1,\ldots,x_k} = \widehat{\mu_{\downarrow i_1,\ldots,i_k}}$ with $i_j = \lceil nx_j\rceil$. As before, we let $(\boldsymbol{x}_i)_{i\ge1}$ be a sequence of independent uniform samples from $[0,1)$.

Corollary 2.7. *If $(\mu_n)_n \overset{D\square}{\to}\mu\in M_\Omega$, then for any integer $k\ge1$ we have $\lim_{n\to\infty}\mathrm{E}\left\|\mu_{n\downarrow\boldsymbol{x}_1,\ldots,\boldsymbol{x}_k} - \mu_{\downarrow\boldsymbol{x}_1,\ldots,\boldsymbol{x}_k}\right\|_1 = 0$.*

Proof. By [22, Theorem 8.6] we can turn μ,μ_n into graphons $w,w_n:[0,1)^2\to\mathscr{P}(\Omega)$ such that for all n

$$\mu = \int_0^1 \delta_{w(\cdot,y)}\,dy,\quad \mu_n = \int_0^1\delta_{w_n(\cdot,y)}\,dy \quad\text{and}\quad D_\square(\mu,\mu_n) = \sup_{U,V\subset[0,1)}\left\|\int_{U\times V} w(x,y)-w_n(x,y)\,dx\,dy\right\|_1.$$

Let $(\boldsymbol{y}_j)_{j\ge1}$ be independent and uniform on $[0,1)$ and independent of $(\boldsymbol{x}_i)_{i\ge1}$. By [29, Theorem 10.7], we have $\lim_{n\to\infty}D_\square(\mu_n,\mu)=0$ iff

$$\lim_{r\to\infty}\limsup_{n\to\infty}\mathrm{E}\left[\max_{I,J\subset[r]}\left\|\sum_{(i,j)\in I\times J} w(\boldsymbol{x}_i,\boldsymbol{y}_j)-w_n(\boldsymbol{x}_i,\boldsymbol{y}_j)\right\|_1\right] = 0. \tag{2.6}$$

Hence, we are left to show that (2.6) implies

$$\forall k\ge1:\lim_{n\to\infty}\mathrm{E}\left\|\mu_{n\downarrow\boldsymbol{x}_1,\ldots,\boldsymbol{x}_k}-\mu_{\downarrow\boldsymbol{x}_1,\ldots,\boldsymbol{x}_k}\right\|_1 = 0. \tag{2.7}$$

To this end, we note that by the strong law of large numbers uniformly for all $x_1,\ldots,x_k\in[0,1]$ and n,

$$\frac1r\sum_{j=1}^r(w(x_1,\boldsymbol{y}_j),\ldots,w(x_k,\boldsymbol{y}_j)) \overset{r\to\infty}{\to} \mu_{\downarrow x_1,\ldots,x_k} \qquad \text{in probability,} \tag{2.8}$$

$$\frac1r\sum_{j=1}^r(w_n(x_1,\boldsymbol{y}_j),\ldots,w_n(x_k,\boldsymbol{y}_j)) \overset{r\to\infty}{\to} \mu_{n\downarrow x_1,\ldots,x_k} \qquad \text{in probability.} \tag{2.9}$$

Hence, if (2.6) holds, then (2.7) follows from (2.8)–(2.9). □

As an application of Corollary 2.7 we obtain

Corollary 2.8. *Assume that $(\mu_n)_n$ is a sequence such that $\mu_n\overset{D\square}{\to}\mu\in M_\Omega$. The following statements are equivalent.*

(i) There is $\sigma\in\Sigma_\Omega$ such that $\mu=\delta_\sigma$.
(ii) For any integer $k\ge2$ we have

$$\lim_{n\to\infty}\mathrm{E}\left\|\mu_{n\downarrow\boldsymbol{x}_1,\ldots,\boldsymbol{x}_k}-\mu_{n\downarrow\boldsymbol{x}_1}\otimes\cdots\otimes\mu_{n\downarrow\boldsymbol{x}_k}\right\|_1 = 0. \tag{2.10}$$

(iii) The condition (2.10) holds for $k=2$.

Proof. The implication (i)⇒(ii) follows from Corollary 2.7 and the step from (ii) to (iii) is immediate. Hence, assume that (iii) holds. Then by Corollary 2.7 and the continuity of the ⊗-operator,

$$\mathrm{E}\left\|\mu_{\downarrow \boldsymbol{x}_1,\boldsymbol{x}_2} - \mu_{\downarrow \boldsymbol{x}_1} \otimes \mu_{\downarrow \boldsymbol{x}_2}\right\|_1 = \lim_{n\to\infty} \mathrm{E}\left\|\mu_{n\downarrow \boldsymbol{x}_1,\boldsymbol{x}_2} - \mu_{n\downarrow \boldsymbol{x}_1} \otimes \mu_{n\downarrow \boldsymbol{x}_2}\right\|_1 = 0. \tag{2.11}$$

Define $\tilde{\sigma} : [0,1) \to \mathscr{P}(\Omega)$ by $x \mapsto \mu_{\downarrow x}$ and assume that $\mu \neq \delta_{\tilde{\sigma}}$. Then $D_\square(\mu, \delta_{\tilde{\sigma}}) > 0$ (by Fact 2.1), whence there exist $B \subset \Sigma_\Omega$, $U \subset [0,1)$, $\omega \in \Omega$ such that

$$\int_B \left[\int_U \sigma_x(\omega) - \tilde{\sigma}_x(\omega)\, dx\right]^2 d\mu(\sigma) > 0. \tag{2.12}$$

However, (2.11) entails

$$\begin{aligned}\int_{\Sigma_\Omega} \left[\int_U \sigma_x(\omega) - \tilde{\sigma}_x(\omega)\, dx\right]^2 d\mu(\sigma) &= \int_{\Sigma_\Omega} \int_U \int_U \sigma_x(\omega)\sigma_y(\omega) - \tilde{\sigma}_x(\omega)\tilde{\sigma}_y(\omega)\, dx\, dy\, d\mu(\sigma) \\ &= \mathrm{E}[\mu_{\downarrow \boldsymbol{x}_1,\boldsymbol{x}_2} - \mu_{\downarrow \boldsymbol{x}_1} \otimes \mu_{\downarrow \boldsymbol{x}_2} | \boldsymbol{x}_1, \boldsymbol{x}_2 \in U] = 0,\end{aligned}$$

in contradiction to (2.12). □

Remark 2.9. *Strictly speaking, the results from* [19, 29] *are stated for graphons with values in* $[0,1]$*, i.e.,* $\mathscr{P}(\Omega)$ *for* $|\Omega| = 2$*. However, they extend to* $|\Omega| > 2$ *directly. For instance, the compactness proof* [29, Chapter 9] *is by way of the regularity lemma, which we extend in Section 2.4 explicitly. Moreover, the sampling result for Corollary 2.7 follows from* [29, Chapter 10] *by viewing* $w : [0,1)^2 \to \mathscr{P}(\Omega)$ *as a family* $(w_\omega)_{\omega\in\Omega}$*,* $w_\omega : (x,y) \mapsto w_{x,y}(\omega) \in [0,1]$*. Finally, the proof of Corollary 2.6 in* [19] *by counting homomorphisms, extends to* $\mathscr{P}(\Omega)$*-valued graphons* [29, Section 17.1].

2.3. **Algebraic properties.** The cut metric is compatible with basic algebraic operations on measures. The following is immediate.

Fact 2.10. *If* $\mu_n \overset{D_\square}{\to} \mu$*,* $\nu_n \overset{D_\square}{\to} \nu$*, then* $\alpha\mu_n + (1-\alpha)\nu_n \overset{D_\square}{\to} \alpha\mu + (1-\alpha)\nu$ *for any* $\alpha \in (0,1)$*.*

The construction of a "product measure" is slightly more interesting. Let Ω, Ω' be finite sets. For $\sigma \in \Sigma_\Omega, \tau \in \Sigma_{\Omega'}$ we define $\sigma \times \tau \in \Sigma_{\Omega\times\Omega'}$ by letting $\sigma \times \tau(x) = \sigma(x) \otimes \tau(x)$, where $\sigma(x) \otimes \tau(x) \in \mathscr{P}(\Omega \times \Omega')$ is the usual product measure of $\sigma(x), \tau(x)$. Further, for $\mu \in M_\Omega, \nu \in M_{\Omega'}$ we define $\mu \times \nu \in M_{\Omega\times\Omega'}$ by

$$\mu \times \nu = \int_{\Sigma_\Omega \times \Sigma_{\Omega'}} \delta_{\sigma\times\tau}\, d\mu \otimes \nu(\sigma, \tau).$$

Clearly, $\mu \times \nu$ is quite different from the usual product measure $\mu \otimes \nu$. However, for *discrete* measures we observe the following.

Fact 2.11. *For* $\mu \in \mathscr{P}(\Omega^n)$ *and* $\nu \in \mathscr{P}(\Omega'^n)$ *we have* $\hat{\mu} \times \hat{\nu} = \widehat{\mu \otimes \nu}$*.*

Proposition 2.12. *If* $\mu_n \overset{D_\square}{\to} \mu \in M_\Omega$*,* $\nu_n \overset{D_\square}{\to} \nu \in M_{\Omega'}$*, then* $\mu_n \times \nu_n \overset{D_\square}{\to} \mu \times \nu$*.*

Proof. Let $\varepsilon > 0$ and choose n_0 large enough so that $D_\square(\mu_n, \mu) < \varepsilon$ and $D_\square(\nu_n, \nu) < \varepsilon$ for all $n > n_0$. By Fact 2.1 there exist couplings γ_n, γ_n' of μ_n, μ and ν_n, ν such that (2.2) is attained. Because $\left\|p \otimes p' - q \otimes q'\right\|_1 \leq \left\|p - q\right\|_1 + \left\|q - q'\right\|_1$ for any $p, q \in \mathscr{P}(\Omega)$, $p', q' \in \mathscr{P}(\Omega)$, we obtain for any $U \subset [0,1)$, $B \subset M_\Omega$, $B' \subset M_\Omega$

$$\left\|\int_{B\times B'} \int_U \sigma \times \sigma'(x) - \tau \times \tau'(x)\, dx\, d\gamma_n \otimes \gamma_n'(\sigma, \tau, \sigma', \tau')\right\|_1 < 2\varepsilon,$$

as desired. □

2.4. **Regularity.** For $\sigma \in \Sigma_\Omega$ and $U \subset [0,1)$ measurable we write

$$\sigma[\omega|U] = \int_U \sigma_x(\omega)\, dx.$$

Moreover, for $\mu \in M_\Omega$ and a measurable $S \subset \Sigma_\Omega$ with $\mu(S) > 0$ we let $\mu[\cdot|S] \in M_\Omega$ be the conditional distribution. Further, let $\boldsymbol{V} = (V_1, \ldots, V_K)$ be a partition of $[0,1)$ into a finite number of pairwise disjoint measurable sets. Similarly, let $\boldsymbol{S} = (S_1, \ldots, S_L)$ be a partition of Σ_Ω into pairwise disjoint measurable sets. We write $\#\boldsymbol{V}, \#\boldsymbol{S}$ for the number K, L of classes, respectively. A measure $\mu \in M_\Omega$ is *ε-regular* with respect to $(\boldsymbol{V}, \boldsymbol{S})$ if there exists $R \subset [\#\boldsymbol{V}] \times [\#\boldsymbol{S}]$ such that the following conditions hold.

REG1: $\lambda(V_i) > 0$ and $\mu(S_j) > 0$ for all $(i,j) \in R$.
REG2: $\sum_{(i,j)\in R} \lambda(V_i)\mu(S_j) > 1 - \varepsilon$.

REG3: for all $(i,j)\in R$ and all $\sigma,\sigma'\in S_j$ we have $\left\|\sigma[\cdot|V_i]-\sigma'[\cdot|V_i]\right\|_1<\varepsilon$.
REG4: if $(i,j)\in R$, then for every $U\subset V_i$ with $\lambda(U)\geq\varepsilon\lambda(V_i)$ and every $T\subset S_j$ with $\mu(T)\geq\varepsilon\mu(S_j)$ we have

$$\left\|\langle\boldsymbol{\sigma}[\cdot|U]\rangle_{\mu[\cdot|T]}-\langle\boldsymbol{\sigma}[\cdot|V_i]\rangle_{\mu[\cdot|S_j]}\right\|_1<\varepsilon.$$

Thus, R is a set of index pairs (i,j) of "good squares" $V_i\times S_j$. **REG1** provides that every good square has positive measure and **REG2** that the total probability mass of good squares is at least $1-\varepsilon$. Further, by **REG3** the averages $\sigma[\cdot|V_i],\sigma'[\cdot|V_i]\in\mathscr{P}(\Omega)$ over V_i of any two $\sigma,\sigma'\in S_j$ are close. Finally, and most importantly, **REG4** requires that the average $\langle\boldsymbol{\sigma}[\cdot|U]\rangle_{\mu[\cdot|T]}$ over a "biggish" sub-square $U\times T$ is close to the mean over the entire square $V_i\times S_j$.

A *refinement* of a partition $(\boldsymbol{V},\boldsymbol{S})$ is a partition $(\boldsymbol{V}',\boldsymbol{S}')$ such that for every pair $(i',j')\in[\#\boldsymbol{V}']\times[\boldsymbol{S}']$ there is a pair $(i,j)\in[\#\boldsymbol{V}]\times[\boldsymbol{S}]$ such that $(V'_{i'},S'_{j'})\subset(V_i,S_j)$.

Theorem 2.13. *For any $\varepsilon>0$ there exists $N=N(\varepsilon,\Omega)$ such that for every $\mu\in M_\Omega$ the following is true. Every partition $(\boldsymbol{V}_0,\boldsymbol{S}_0)$ with $\#\boldsymbol{V}_0+\#\boldsymbol{S}_0\leq 1/\varepsilon$ has a refinement $(\boldsymbol{V},\boldsymbol{S})$ such that $\#\boldsymbol{V}+\#\boldsymbol{S}\leq N$ with respect to which μ is ε-regular.*

In light of Proposition 2.4, Theorem 2.13 would follow from the regularity lemma for graphons [29, Lemma 9.16] if we were to drop condition **REG3**. In fact, adapting the standard proof from [36] to accommodate **REG3** is not difficult. For the sake of completeness we carry this out in detail in Section 2.5.

A regularity lemma for measures on Ω^n was proved in [8]. But even in the discrete case Theorem 2.13 gives a stronger result. The improvement is that **REG4** above holds for all "small sub-squares" $U\times T$ simultaneously.

How does the concept of regularity connect with the cut metric? For a partition $\boldsymbol{V}$ of $[0,1]$ and $\sigma\in\Sigma_\Omega$ define $\sigma[\cdot|\boldsymbol{V}]\in\mathscr{W}_\Omega$ by

$$\sigma_x[\omega|\boldsymbol{V}]=\sum_{i\in[\#\boldsymbol{V}]}\mathbf{1}\{x\in V_i\}\sigma_x[\omega|V_i].$$

Thus, $\sigma[\cdot|\boldsymbol{V}]:[0,1)\to\mathscr{P}(\Omega)$ is constant on the classes of $\boldsymbol{V}$. Further, for a pair $(\boldsymbol{V},\boldsymbol{S})$ of partitions and $\mu\in M_\Omega$ let

$$\mu[\cdot|\boldsymbol{V},\boldsymbol{S}]=\sum_{i\in[\#\boldsymbol{S}]}\delta_{\int_{S_i}\sigma[\cdot|\boldsymbol{V}]d\mu(\sigma)}.$$

Hence, $\mu[\cdot|\boldsymbol{V},\boldsymbol{S}]\in M_\Omega$ is supported on a discrete set of functions $[0,1)\to\mathscr{P}(\Omega)$ that are constant on the classes of $\boldsymbol{V}$. We might think of $\mu[\cdot|\boldsymbol{V},\boldsymbol{S}]$ as the "conditional expectation" of μ with respect to $(\boldsymbol{V},\boldsymbol{S})$.

Proposition 2.14. *Let $\varepsilon>0$ and assume that μ is ε-regular w.r.t. $(\boldsymbol{V},\boldsymbol{S})$. Then $D_\square(\mu,\mu[\cdot|\boldsymbol{V},\boldsymbol{S}])<2\varepsilon$.*

Proof. Let $\sigma^{(i)}=\int_{S_i}\sigma[\cdot|\boldsymbol{V}]d\mu(\sigma)$. We define a coupling γ of $\mu,\mu[\cdot|\boldsymbol{V},\boldsymbol{S}]$ in the obvious way: for a measurable $X\subset S_i$ let $\gamma(X\times\{\sigma^{(i)}\})=\mu(X)$. Now, let $U\subset[0,1]$ and $B\subset\Sigma_\Omega^2$ be measurable. Due to the construction of our coupling we may assume that $B=\bigcup_i B_i\times\{\sigma^{(i)}\}$ for certain sets $B_i\subset S_i$. Moreover, let $U_j=U\cap V_j$. Then

$$\left\|\int_B\int_U\sigma(x)-\tau(x)dxd\eta(\sigma,\tau)\right\|_1\leq\sum_{(i,j):\mu(S_i)\lambda(V_j)>0}\mu(S_i)\lambda(V_j)\left\|\int_{B_i}\int_{U_j}\sigma(x)-s_i(x)\frac{dx}{\lambda(V_j)}\frac{d\mu(\sigma)}{\mu(S_i)}\right\|_1.$$

By **REG1** and **REG4** the last expression is less than 2ε. □

Corollary 2.15. *For any $\varepsilon>0$ there exists $N=N(\varepsilon)>0$ such that for any $\mu\in M_\Omega$ there exist $\sigma_1,\ldots,\sigma_N\in\Sigma_\Omega$ and $w=(w_1,\ldots,w_N)\in\mathscr{P}([N])$ such that $D_\square\left(\mu,\sum_{i=1}^k w_i\delta_{\sigma_i}\right)<\varepsilon$.*

Proof. This is immediate from Theorem 2.13 and Proposition 2.14. □

2.5. **Proof of Theorem 2.13.** Following the path beaten in [8, 36, 37], we define the index of $(\boldsymbol{V},\boldsymbol{S})$ as

$$\mathrm{ind}_\mu(\boldsymbol{V},\boldsymbol{S})=\mathrm{E}\left\langle\mathrm{Var}[\boldsymbol{\sigma}_{\boldsymbol{x}}[\omega]|\boldsymbol{V},\boldsymbol{S}]\right\rangle_\mu=\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\sum_{i=1}^{\#\boldsymbol{V}}\sum_{j=1}^{\#S_j}\int_{S_j}\int_{V_i}\left(\sigma_x(\omega)-\int_{S_j}\int_{V_i}\sigma_y(\omega)\frac{dy}{\lambda(V_i)}\frac{d\mu(\sigma)}{\mu(S_j)}\right)^2dxd\mu(\sigma).$$

There is only one simple step that we add to the proof from [36]. Namely, following [8], we begin by refining the partition $\boldsymbol{S}_0$ to guarantee **REG3**. Specifically, the compact set $\mathscr{P}(\Omega)$ has a partition into a finite number of sets $\boldsymbol{Q}=(Q_1,\ldots,Q_K)$ such that $\left\|\mu-\mu'\right\|_1\leq\varepsilon$ for all $\mu,\mu'\in Q_i$, $i\in[K]$. Now, let $\boldsymbol{V}(1)=\boldsymbol{V}_0$ and let $\boldsymbol{S}(1)$ be the coarsest refinement of $\boldsymbol{S}_0$ such that for every $i\in[\#\boldsymbol{V}_1]$, $j\in\#\boldsymbol{S}_1$ there is $k\in[K]$ such that $\sigma[\cdot|V_i]\in Q_k$ for all $\sigma\in S_j(1)$. Then $\#\boldsymbol{S}(1)\leq K\#\boldsymbol{S}_0$.

Starting from $(\boldsymbol{V}(1),\boldsymbol{S}(1))$, we construct a sequence $(\boldsymbol{V}(t),\boldsymbol{S}(t))$ of partitions inductively. The construction stops once μ is ε-regular w.r.t. $(\boldsymbol{V}(t),\boldsymbol{S}(t))$, in which case we are done. Assuming otherwise, consider the set $\bar{R}(t)\subset[\#\boldsymbol{V}(t)]\times[\#\boldsymbol{S}(t)]$ of (i,j) such that **REG4** fails to hold on (V_i,S_j). Then

$$\sum_{(i,j)\in\bar{R}(t)}\lambda(V_i)\mu(S_j)\geq\varepsilon. \tag{2.13}$$

Further, for each $(i,j)\in\bar{R}$ there exist $U_i\subset V_i$, $\lambda(U_i)\geq\varepsilon\lambda(V_i)$, $T_j\subset S_j$, $\mu(T_j)\geq\varepsilon\mu(S_j)$ and ω_{ij} such that

$$\Big|\big\langle\boldsymbol{\sigma}[\omega_{ij}|U_i]\big\rangle_{\mu[\cdot|T_j]}-\big\langle\boldsymbol{\sigma}[\omega_{ij}|V_i]\big\rangle_{\mu[\cdot|S_j]}\Big|\geq\frac{\varepsilon}{|\Omega|}. \tag{2.14}$$

Obtain the partition $(\boldsymbol{V}(t,i,j),\boldsymbol{S}(t,i,j))$ from $(\boldsymbol{V}(t),\boldsymbol{S}(t))$ by splitting V_i and S_j into the sub-classes U_i, $V_i\setminus U_i$ and S_j, $S_j\setminus T_j$. Clearly, $\#\boldsymbol{V}_{ij}\leq2\#\boldsymbol{V}$ and $\#\boldsymbol{S}_{ij}\leq2\#\boldsymbol{S}$, respectively. Then

$$\begin{aligned}\mathrm{E}[\big\langle\mathrm{Var}[\boldsymbol{\sigma}_{\boldsymbol{x}}(\omega_{ij})|\boldsymbol{V}(t),\boldsymbol{S}(t)]\big\rangle_{\mu[\cdot|S_j]}|V_i]&=\mathrm{E}[\big\langle\mathrm{Var}[\boldsymbol{\sigma}_{\boldsymbol{x}}(\omega_{ij})|\boldsymbol{V}(t,i,j),\boldsymbol{S}(t,i,j)]\big\rangle_{\mu[\cdot|S_j]}|V_i]\\&\quad+\mathrm{E}[\big\langle\mathrm{Var}[\mathrm{E}[\boldsymbol{\sigma}_{\boldsymbol{x}}(\omega_{ij})|\boldsymbol{V}(t,i,j),\boldsymbol{S}(t,i,j)]|\boldsymbol{V}(t),\boldsymbol{S}(t)]\big\rangle_{\mu[\cdot|S_j]}|V_i].\end{aligned} \tag{2.15}$$

Moreover, (2.14) implies that on $V_i\times S_j$ we have

$$\mathrm{Var}[\mathrm{E}[\boldsymbol{\sigma}_{\boldsymbol{x}}[\omega]|\boldsymbol{V}(t,i,j),\boldsymbol{S}(t,i,j)]|\boldsymbol{V}(t),\boldsymbol{S}(t)]\geq\frac{\lambda(U_i)}{\lambda(V_i)}\frac{\mu(T_j)}{\mu(S_j)}\Big|\big\langle\boldsymbol{\sigma}[\omega_{ij}|U_i]\big\rangle_{\mu[\cdot|T_j]}-\big\langle\boldsymbol{\sigma}[\omega_{ij}|V_i]\big\rangle_{\mu[\cdot|S_j]}\Big|^2. \tag{2.16}$$

Combining (2.14), (2.15) and (2.16), we obtain

$$\mathrm{E}[\big\langle\mathrm{Var}[\boldsymbol{\sigma}_{\boldsymbol{x}}(\omega_{ij})|\boldsymbol{V}(t),\boldsymbol{S}(t)]\big\rangle_{\mu[\cdot|S_j]}|V_i]\geq\mathrm{E}[\big\langle\mathrm{Var}[\boldsymbol{\sigma}_{\boldsymbol{x}}(\omega_{ij})|\boldsymbol{V}(t,i,j),\boldsymbol{S}(t,i,j)]\big\rangle_{\mu[\cdot|S_j]}|V_i]+\lambda(U_i)\mu(T_j)\frac{\varepsilon^2}{|\Omega|^3}. \tag{2.17}$$

Now, let $(\boldsymbol{V}(t+1),\boldsymbol{S}'(t+1))$ denote the coarsest common refinement of all the partitions $(\boldsymbol{V}(t,i,j),\boldsymbol{S}(t,i,j))_{i,j\in\bar{R}}$. Finally, obtain $\boldsymbol{S}(t+1)$ from $\boldsymbol{S}'(t+1)$ as in the very first step by splitting each class $\boldsymbol{S}'_j(t+1)$ into classes $(\boldsymbol{S}'_{j,k}(t+1))_{k\leq K}$ such that $\big\|\sigma[\cdot|\boldsymbol{V}_i(t+1)]-\sigma'[\cdot|\boldsymbol{V}_i(t+1)]\big\|_1<\varepsilon$ for all $i\in[\#\boldsymbol{V}(t+1)]$. Then (2.17) and the monotonicity of the conditional variance imply

$$\mathrm{ind}_\mu(\boldsymbol{V}(t+1),\boldsymbol{S}(t+1))\leq\mathrm{ind}_\mu(\boldsymbol{V}(t),\boldsymbol{S}(t))-\frac{\varepsilon^2}{|\Omega|^3}\sum_{(i,j)\in\bar{R}}\lambda(U_i)\mu(T_j)\leq\mathrm{ind}_\mu(\boldsymbol{V}(t),\boldsymbol{S}(t))-\frac{\varepsilon^5}{|\Omega|^3}. \tag{2.18}$$

Since the index lies between 0 and 1, (2.18) implies that the construction stops after at most $\varepsilon^{-5}|\Omega|^3$ steps.

3. The replica symmetric solution

In this section we apply the notion of convergence and the results from Section 2 to Gibbs measures induced by random graphs. Many of the arguments build upon the work of Panchenko on asymptotic Gibbs measures [33].

3.1. **Random factor graphs.**

A remarkably wide variety of problems in combinatorics can be described in terms of *factor graphs*. These are bipartite graphs with two types of vertices called *variable nodes* and *constraint nodes*. The variable nodes can be assigned "spins" from a finite set Ω and each constraint node is decorated with a weight function that assigns every spin configuration of its adjacent variable nodes a positive weight. Natural examples of such models occur in combinatorics, mathematical physics or information theory [31, 35]. We shall see a few concrete examples in just a moment.

Let us first attempt an abstract, fairly comprehensive definition (see [8] for an even more general setup). Suppose that Ω is a given finite set of *spins* and that Ψ is a set of functions $\psi:\Omega^{k_\psi}\to(0,\infty)$ of arity $k_\psi\geq1$. A (Ω,Ψ)-*factor graph* $G=(V_G,F_G,(\psi_a)_{a\in F_G},\partial_G)$ consists of

- **GR1:** (finite or) countable disjoint sets V_G,F_G,
- **GR2:** a map $a\in F_G\mapsto\psi_a\in\Psi$,
- **GR3:** a map $\partial_G:F_G\to\bigcup_{l\geq1}V_G^l$ such that $\partial_Ga=(\partial_G(a,j))_j\in V_G^{k_{\psi a}}$ for all $a\in F_G$ and such that for every $x\in V_G$ the set $\{a\in F_G:\exists j\in[k_{\psi_a}]:x=\partial_G(a,j)\}$ is finite.

Let us introduce the shorthand $k_a = k_{\psi_a}$ for the arity of the constraint a. Moreover, if $\sigma : V_G \to \Omega$, then we let $\sigma(\partial_G a) = (\sigma(\partial_G(a,1)),\dots,\sigma(\partial_G(a,k_a)))$.

A finite factor graph G naturally induces a probability measure on the set Ω^{V_G} of all possible assignments: the *Gibbs measure* of G is defined by

$$\mu_G : \Omega^{V_G} \to (0,1), \quad \sigma \mapsto Z_G^{-1} \prod_{a \in F_G} \psi_a(\sigma(\partial_G a)), \qquad \text{where} \qquad Z_G = \sum_{\sigma \in \Omega^{V_G}} \prod_{a \in F_G} \psi_a(\sigma(\partial_G a)) \tag{3.1}$$

is the *partition function* of G. Thus, the probability mass that μ_G assigns to σ is proportional to the weights $\psi_a(\sigma(\partial_G a))$ that the constraint nodes a assign to the spin configurations of their incident variables.

Naturally, we can view a factor graph as a bipartite graph with node sets V_G and F_G such that each $a \in F_G$ is adjacent to the variable nodes $\partial_G(a,j)$ for $j \le k_{\psi_a}$. Hence, we call them the *neighbors* of a and take license to just write $\partial_G a$ for the set of neighbors. Conversely, we write $\partial_G x$ for the set of constraint nodes a such that $x \in \partial_G a$. By **GR3** $\partial_G x$ is a finite set for every x. Further, a *rooted factor graph* is a connected factor graph together with a distinguished variable node r, its *root*.

However, we keep in mind that the "bipartite graph" point of view looses a bit of information. Indeed, the neighbors of a are *ordered*. This is important in evaluating (3.1) because the ψ_a need not be permutation invariant.

Example 3.1 (The Ising model on the grid $\mathbb{Z}^2$)**.** *Let G_n be a finite subgraph of $\mathbb{Z}^2$ (say a large box), $\Omega = \{\pm 1\}$, and $\psi : \{\pm 1\}^2 \to (0,\infty)$ be defined by $\psi(x_1,x_2) = \exp(\beta \cdot x_1 x_2)$ for a fixed real number β. Then the Ising model on G_n is defined by*

$$\mu_{G_n}(\sigma) = \frac{\prod_{(u,v)\in E(G_n)} \psi(\sigma(u),\sigma(v))}{Z_{G_n}},$$

where

$$Z_{G_n} = \sum_{\sigma \in \{\pm 1\}^{V(G)}} \prod_{(u,v)\in E(G_n)} \psi(\sigma(u),\sigma(v)).$$

If $\beta > 0$ we say the model is ferromagnetic (like spins preferred across edges) and if $\beta < 0$, the model is antiferromagnetic.

Example 3.2 (The (positive-temperature) k-SAT model)**.** *The k-SAT model is an example of a factor model with multiple constraint types. We take $\Omega = \{\pm 1\}$ and have 2^k constraint types, each of arity k, indexed by vectors $\boldsymbol{c} \in \{\pm 1\}^k$, with*

$$\psi^{(\boldsymbol{c})}(\boldsymbol{x}) = \exp(-\beta \mathbf{1}\{\boldsymbol{c} \cdot \boldsymbol{x} = -k\})$$

for $\boldsymbol{x} \in \{\pm 1\}^k$. The temperature parameter β is a fixed real number controlling how much satisfied clauses are preferred to unsatisfied clause. The 2^k different constraint function types correspond to the 2^k different ways to assign signs to k variables that appear together in a k-CNF clause. We have $\boldsymbol{c} \cdot \boldsymbol{x} = -k$ if all k signed variables are -1, in which case the clause is unsatisfied. We form a random instance of the k-SAT model by choosing a random number of constraints of each of the 2^k types to form F_G, and to each constraint $a \in F_G$ attaching a uniformly random ordered set of k variable nodes $x_{a_1},\dots x_{a_k}$ from $V_G = \{x_1,\dots x_n\}$. The Gibbs measure is then a probability distribution over all assignments to the n variables given by

$$\mu_G(\sigma) = \frac{\prod_{a\in F_G} \psi_a(\sigma(x_{a_1}),\dots \sigma(x_{a_k}))}{Z_G} \qquad \textit{where} \qquad Z_G = \sum_{\sigma} \prod_{a \in F_G} \psi_a(\sigma(x_{a_1}),\dots \sigma(x_{a_k})).$$

Studying μ_G can be viewed as a generalization of the MAX k-SAT problem [3]*. In fact, the maximum number of clauses that can be satisfied simultaneously comes out as $|F_G| + \lim_{\beta\to\infty} \frac{\partial}{\partial \beta} \ln Z_G$.*

Apart from natural geometric factor graph models such as the Ising model, there is substantial interest in models where the geometry of interactions is random like in Example 3.2. For instance, such models appear in the statistical mechanics of disordered systems, coding theory and, of course, the theory of random graphs itself [23, 31, 35]. Perhaps the simplest and most natural way of defining such models is by extension of the Erdős-Rényi model. Hence, given a set Ψ of possible weight functions and a sequence $\rho = (\rho_\psi)_{\psi\in\Psi}$ of positive reals we define the random factor graph $\boldsymbol{G}_n = \boldsymbol{G}_n(\Psi,\rho)$ as follows. The set of variable nodes is $V_n = \{x_1,\dots,x_n\}$. Moreover, choosing $m_\psi = \mathrm{Po}(n\rho_\psi)$ independently for each $\psi \in \Psi$, we define the set of constraint nodes as

$$F_n = \{a_{\psi,i} : i \le m_\psi \text{ for all } \psi \in \Psi\}.$$

Further, independently for each $a_{\psi,i} \in F_n$ choose $\partial_{\boldsymbol{G}_n} a_{\psi,i} \in V_n^{k_\psi}$ uniformly at random.

A random factor graph $\boldsymbol{G}_n$ induces a Gibbs measure $\mu_{\boldsymbol{G}_n}$ on Ω^n via (3.1). Thus, we could just apply the "limit theory" for discrete measures from Section 2 to the sequence $(\mu_{\boldsymbol{G}_n})_n$. Indeed, this is essentially what Panchenko [33] does (using the Aldous-Hoover representation instead of Section 2). However, the geometry of the sparse random graph $\boldsymbol{G}_n$ contains additional information that is not directly encoded in the measure $\mu_{\boldsymbol{G}_n}$. Hence, the basic idea in the following is to study the convergence of the sequence $(\mu_{\boldsymbol{G}_n})_n$ of measures jointly with the convergence of the geometry of the factor graph $\boldsymbol{G}_n$. This will enable us to identify the limit of $(\mu_{\boldsymbol{G}_n})_n$ with a "geometric" measure on a (possible infinite) random tree. In particular, we aim to use these insights to get a handle on the *free energy* of the model, defined as

$$\lim_{n\to\infty} \frac{1}{n} \mathrm{E}[\ln Z_{\boldsymbol{G}_n}], \tag{3.2}$$

provided that the limit even exists. (Of course, Ψ, ρ remain fixed as n grows.)

A similar "geometric" approach was pursued in [8] for a more general class of factor graph models, and without a proper notion of convergence of the sequence of Gibbs measures. For the more specific models studied here we will obtain somewhat stronger results from simpler proofs.

3.2. **Local weak convergence.** To carry out this program we need to set up a notion of a "limit" of the geometry of $\boldsymbol{G}_n$ as $n \to \infty$. To adapt the appropriate formalism of "local weak convergence" [6, 12, 29] to our context, let $G = (V_G, F_G, \partial_G, (\psi_a)_{a\in F_G}, r)$, $G' = (V_{G'}, F_{G'}, \partial_{G'}, (\psi_a)_{a\in F_G}, r')$ be two rooted factor graphs. An *isomorphism* $f : G \to G'$ is a bijection $f : V_T \cup F_T \to V_{T'} \cup F_{T'}$ such that

ISM1: $f(V_T) = V_{T'}$, $f(F_T) = F_{T'}$, $f(r) = r'$,
ISM2: $\psi_{f(a)} = \psi_a$ for all $a \in F_T$,
ISM3: $\partial_{T'}(f(a), j) = f(\partial_T(a, j))$ for all $j \in [k_{\psi_a}]$.

We write $G \cong G'$ if there is an isomorphism $G \to G'$.

Let $[G]$ be the isomorphism class of G and let $\mathfrak{G}$ be the set of all isomorphism classes. Further, for an integer $\ell \geq 1$ let $\partial^\ell G$ be obtained from G by deleting all (variable and constraint) nodes whose distance from the root exceeds 2ℓ. Then it makes sense to write $\partial^\ell [G]$, because $G \cong H$ implies that $\partial^\ell G \cong \partial^\ell H$ for all ℓ.

To be allowed to use standard graph terminology for isomorphism classes, let us pick one representative G_0 of every isomorphism class $[G]$ arbitrarily. Hence, if we speak, e.g., of "the neighbor of the root of $[G]$", we refer to the corresponding object in the chosen representative G_0.

We endow $\mathfrak{G}$ with the coarsest topology that makes all the functions

$$\mathfrak{G} \to \{0,1\}, \qquad [G] \mapsto \mathbf{1}\{\partial^\ell G \cong \partial^\ell H\} \qquad ([H] \in \mathfrak{G}) \tag{3.3}$$

continuous. Moreover, let $\mathfrak{T} \subset \mathfrak{G}$ be the set of all isomorphism classes of all acyclic rooted factor graphs with the induced topology. The spaces $\mathfrak{G}, \mathfrak{T}$ are Polish [6, 12]. Hence, so are the spaces $\mathscr{P}(\mathfrak{G})$, $\mathscr{P}(\mathfrak{T})$ of probability measures on $\mathfrak{G}, \mathfrak{T}$ with the weak topology. Additionally, we equip the spaces $\mathscr{P}^2(\mathfrak{G})$, $\mathscr{P}^2(\mathfrak{T})$ of probability measures on $\mathscr{P}(\mathfrak{G})$, $\mathscr{P}(\mathfrak{T})$ with the weak topology as well.

For a factor graph G and a variable node v let $G_{\uparrow v}$ be the connected component of v in G rooted at v. Thus, $G_{\uparrow v}$ is a rooted factor graph. Similarly, if (G, r) is a rooted factor graph, then $G_{\uparrow v} = (G, v)$ is obtained by re-rooting at v. Each factor graph G induces an empirical distribution on $\mathfrak{G}$, namely

$$\Lambda_G = |V_G|^{-1} \sum_{x\in V_G} \delta_{G_{\uparrow x}} \in \mathscr{P}(\mathfrak{G}).$$

Hence, the random factor graph gives rise to a distribution

$$\Lambda_n = \mathrm{E}[\delta_{\Lambda_{\boldsymbol{G}_n}}] \in \mathscr{P}^2(\mathfrak{G}).$$

Due to the definition (3.3) of the topology, this measure captures the distribution of the "local structure" of $\boldsymbol{G}_n$, i.e., the "statistics" of the bounded-size neighborhoods.

Guided by the example of the Erdős-Rényi random graph, we expect that the local structure of $\boldsymbol{G}_n$ is described by a branching process. Specifically, starting from a single variable node $\mathcal{V}_0 = \{x_0\}$ and with T_0 the tree consisting of x_0 only, we build a sequence of random trees $(\boldsymbol{T}_\ell)_\ell$ as follows. Let $(A_{\psi,x})_{\psi\in\Psi, x\in\mathcal{V}_\ell}$ be a family of independent random variables such that $A_{\psi,x}$ has distribution $\mathrm{Po}(\rho_\psi)$.

Now, obtain $\boldsymbol{T}_{\ell+1}$ from $\boldsymbol{T}_\ell$ by attaching $Y_{\psi,x}$ children, which are constraint nodes with weight function ψ, to each $x \in \mathcal{V}_\ell$. For each of them independently choose the position that the parent variable occupies uniformly and

independently from $[k_\psi]$ and attach $k_\psi - 1$ further variable nodes. Finally, let $\mathcal{V}_{\ell+1}$ be the set of variable nodes of $\boldsymbol{T}_{\ell+1}$ at distance precisely $2(\ell+1)$ from the root.

Let $\vartheta_\ell \in \mathscr{P}(\mathfrak{T}_\ell)$ be the distribution of the random tree $\boldsymbol{T}_\ell$. Because the topology is generated by the functions (3.3), $(\vartheta_\ell)_{\ell\geq 1}$ is a Cauchy sequence. Since $\mathscr{P}(\mathfrak{T})$ is a complete space, there exists a limit $\vartheta \in \mathscr{P}(\mathfrak{T})$. Further, write $\boldsymbol{T}$ for a random (possibly infinite) tree drawn from ϑ.

Proposition 3.3. *We have* $\lim_{n\to\infty} \Lambda_n = \delta_\vartheta$.

Proof. For $T \in \mathfrak{T}$ and $\ell \geq 1$ let $Q_{T,\ell}(\boldsymbol{G}_n)$ be the number of variable nodes x of $\boldsymbol{G}_n$ such that $\partial^\ell \boldsymbol{G}_{n\uparrow x} \cong \partial^\ell T$. Unravelling the construction of the topology via (3.3), we see that $\lim_{n\to\infty} \Lambda_n = \delta_\vartheta$ iff $n^{-1}Q_{T,\ell}(\boldsymbol{G}_n)$ converges in probability to $\mathrm{P}\left[\partial^\ell \boldsymbol{T} = \partial^\ell T\right]$ for every T, ℓ.

Hence, fix T, ℓ and assume that n is large. Further, let $\boldsymbol{G}'_n$ be a random factor graph with variable nodes V_n in which for every $\psi \in \Psi$ each of the n^{k_ψ} possible constraint with weight function ψ is present with probability $p_\psi = \rho_\psi n^{1-k_\psi}$ independently. Then the number M_ψ of constraints of type ψ has distribution $\mathrm{Bin}(n^{k_\psi}, p_\psi)$, and they are mutually independent. Because the total variation distance of M_ψ and $\mathrm{Po}(n\rho_\psi)$ is $o(1)$ as $n \to \infty$, the same is true of the random graph distributions $\boldsymbol{G}_n, \boldsymbol{G}'_n$.

We are now going to show by induction on ℓ that there is a coupling of $[\partial^\ell_{\boldsymbol{G}_n} x_1]$ and $\partial^\ell \boldsymbol{T}$ such that both coincide with probability $1 - o(1)$. For $\ell = 0$ there is nothing to show as both graphs consist of the root only. To proceed from ℓ to $\ell+1$, let W_ℓ be the set of variable nodes at distance precisely 2ℓ from x_1 in $\boldsymbol{G}'_n$ and let V_ℓ be the set of variable nodes at distance precisely 2ℓ from $r_{\boldsymbol{T}}$. Further, condition on the event $\mathscr{E}_\ell = \{[\partial^\ell_{\boldsymbol{G}_n} x_1] = \partial^\ell \boldsymbol{T}\}$ and fix an isomorphism

$$\varphi : \partial^\ell_{\boldsymbol{G}_n} x_1 \to \partial^\ell \boldsymbol{T}.$$

Moreover, let $\mathscr{X}$ be the event that the random factor graph $\boldsymbol{G}_n$ either contains a constraint node a such that $|\partial_{\boldsymbol{G}_n} a \cap W_\ell| \geq 2$ or a constraint nodes b, c such that $|\partial_{\boldsymbol{G}_n} b \cap W_\ell|, |\partial_{\boldsymbol{G}_n} c \cap W_\ell| = 1$ and $\partial_{\boldsymbol{G}_n} b \cap \partial_{\boldsymbol{G}_n} c \setminus W_\ell \neq \emptyset$. Because the tree T remains fixed as we let $n \to \infty$, the number of possible a, b, c with these properties is $O(n^{k_\psi - 2})$ for every ψ. Therefore, $\mathrm{P}[\mathscr{X}|\mathscr{E}_\ell] = O(1/n)$. Furthermore, for every $x \in W_\ell$ let number $D_{x,\psi}$ be the number of constraint nodes a of type ψ such that $\partial_{\boldsymbol{G}_n} a \cap \bigcup_{l\leq \ell} W_l = \{x\}$. Given $\mathscr{E}_\ell \cap \mathscr{X}$, $D_{x,\psi}$ has distribution $k_\psi \cdot \mathrm{Bin}((n - |\bigcup_{l\leq\ell} W_l| - O(1))^{k_\psi - 1}, q_\psi)$, and the $D_{x,\psi}$ are asymptotically independent. Analogously, let $d_{x,\psi}$ be the number of constraint nodes a of type ψ that are children $x \in V_\ell$. Then $d_{x,\psi}$ has distribution $\mathrm{Po}(k_\psi \rho_\psi)$ by the construction of $\boldsymbol{T}$, and the $d_{x,\psi}$ are mutually independent. Consequently, since $|W_\ell| = O(1)$ as $n \to \infty$, there exists a coupling of the vectors $(d_{x,\psi})_{x\in W_\ell, \psi\in\Psi}$, $(D_{x,\psi})_{x\in V_\ell, \psi\in\Psi}$ such that $\mathrm{P}\left[\forall x, \psi : d_{x,\psi} = D_{\varphi(x),\psi} | \mathscr{X} \cap \mathscr{E}_\ell\right] = 1 - o(1)$. Hence, we obtain a coupling of $\partial^{\ell+1}_{\boldsymbol{G}_n} x_1$ and $\partial^{\ell+1}\boldsymbol{T}$ such that $\mathrm{P}[[\partial^{\ell+1}_{\boldsymbol{G}_n} x_1] = \partial^{\ell+1}\boldsymbol{T}] = 1 - o(1)$, as desired.

Because the random factor graph model is invariant under permutations of the variable nodes, the existence of this coupling implies that $\mathrm{E}[Q_{T,\ell}(\boldsymbol{G}_n)] \sim n\mathrm{P}\left[\partial^\ell \boldsymbol{T} = \partial^\ell T\right]$. To estimate the second moment $\mathrm{E}[Q_{T,\ell}(\boldsymbol{G}_n)^2]$, we repeat the argument from the previous paragraph for the two variable nodes x_1, x_2 to show that

$$\mathrm{P}\left[\partial^\ell_{\boldsymbol{G}_n} x_1 = T, \partial^\ell_{\boldsymbol{G}_n} x_2 = T'\right] \sim \mathrm{P}\left[\partial^\ell \boldsymbol{T} = \partial^\ell T\right] \mathrm{P}\left[\partial^\ell \boldsymbol{T} = \partial^\ell T\right]. \tag{3.4}$$

Once more by permutation-invariance, (3.4) implies that $\mathrm{E}[Q_{T,\ell}(\boldsymbol{G}_n)^2] \sim \mathrm{E}[Q_{T,\ell}(\boldsymbol{G}_n)]^2$. Finally, the desired convergence in probability follows from Chebyshev's inequality. □

3.3. **Replica symmetry.** Having discussed the meaning of "convergence of the local structure", let us now return to the convergence of the Gibbs measure $\mu_{\boldsymbol{G}_n}$ itself. The "cavity method", a non-rigorous but sophisticated approach from statistical physics [26, 31], predicts a relatively simple formula for the free energy (3.2) if $\mu_{\boldsymbol{G}_n}$ converges to an atom δ_w, $w \in \Sigma_\Omega$, in the metric $\Delta_\square(\cdot,\cdot)$. This convergence assumption roughly coincides with the *replica symmetry* condition from physics [26, 33]. According to the cavity method, in the replica symmetric case the free energy can be calculated by applying an explicit functional, the *Bethe free energy* [38], to a fixed point of a message passing scheme called *Belief Propagation* [31]. We are going to vindicate this prediction.

But before we introduce Belief Propagation and the Bethe free energy, let us briefly discuss the replica symmetry assumption. Formally, we are going to assume that there is a function $w \in \Sigma_\Omega$ such that

$$\lim_{n\to\infty} \mathrm{E}\left[\Delta_\square(\mu_{\boldsymbol{G}_n}, \delta_w)\right] = 0. \tag{3.5}$$

In other words, $\mu_{\boldsymbol{G}_n}$ converges to δ_w in probability with respect to $\Delta_\square(\cdot,\cdot)$.

The assumption (3.5) holds in all examples of random factor graph models where we currently have an at least somewhat explicit formula for the free energy (to our knowledge). For example, this includes all cases in which

the free energy can be computed by the "second moment method" (e.g., [1, 2, 17]). Indeed, in these examples $w : x \in [0,1) \mapsto p \in \mathscr{P}(\Omega)$ is a *constant* function. However, there are replica symmetric models in which the limiting density w is not constant.

Instead of relying on the second moment method, the condition (3.5) can be checked (and the function w can be computed) by studying spatial mixing properties of the Gibbs measure; for an example see [9]. Let us give a simple generic proof that the *non-reconstruction condition*, a spatial mixing property, entails (3.5); this was predicted in [26].

For a random factor graph $\boldsymbol{G}_n$, a variable node $x \in V_n$ and $\ell \geq 1$ let $\nabla_\ell(\boldsymbol{G}_n, x)$ be the σ-algebra on Ω^{V_n} generated by the events $\{\boldsymbol{\sigma}(y) = \omega\}$ for all $\omega \in \Omega$ and all variables y at distance greater than 2ℓ from x. Thus, in the measure $\mu_{\boldsymbol{G}_n}[\cdot|\nabla_\ell(G,x)]$ we condition on all the values of all the variable nodes at distance greater than 2ℓ from x. The random factor graph model has the *non-reconstruction property* if

$$\lim_{\ell\to\infty}\lim_{n\to\infty}\mathrm{E}\left\langle\left\|\mu_{\downarrow x_1} - \mu_{\downarrow x_1}[\cdot|\nabla_\ell(\boldsymbol{G}_n, x_1)]\right\|_1\right\rangle_{\mu_{\boldsymbol{G}_n}} = 0. \tag{3.6}$$

To parse (3.6), we note that the outer expectation $\mathrm{E}[\cdot]$ refers to the choice of the random factor graph $\boldsymbol{G}_n$. Further, the outer mean $\langle\cdot\rangle_{\boldsymbol{G}_n}$ over the Gibbs measure of $\boldsymbol{G}_n$ generates the random boundary condition. We then compare the conditional marginal $\mu_{\downarrow x_1}[\cdot|\nabla_\ell(\boldsymbol{G}_n x_1)]$ given the boundary condition with the unconditional marginal $\mu_{\downarrow x_1}$. Because the distribution of $\boldsymbol{G}_n$ is invariant under permutations of the variables, the choice of the variable x_1 in (3.6) is irrelevant. Hence, (3.6) provides that the impact of a random boundary condition on the marginal of any specific variable x_i diminishes in the limit $\ell, n \to \infty$.

Proposition 3.4. *Assume that* $\lim_{n\to\infty}\mathrm{E}[\Delta_\square(\mu_{\boldsymbol{G}_n}, \mu)] = 0$ *for some* $\mu \in M_\Omega$. *If (3.6) holds, then there exists* $w \in \Sigma_\Omega$ *such that* $\bar{\mu} = \overline{\delta_w}$*; thus, (3.5) holds.*

Proof. We apply an argument from [32] developed for the "stochastic block model" to our setup. Due to Corollary 2.8 and because the distribution of $\boldsymbol{G}_n$ is invariant under permutations of the variables it suffices to prove

$$\lim_{n\to\infty}\mathrm{E}\left\|\mu_{\boldsymbol{G}_n\downarrow x_1,x_2} - \mu_{\boldsymbol{G}_n\downarrow x_1}\otimes\mu_{\boldsymbol{G}_n\downarrow x_2}\right\|_1 = 0. \tag{3.7}$$

Hence, assume that (3.6) holds but (3.7) does not. Then there exist $\omega_1, \omega_2 \in \Omega$ and $0 < \varepsilon < 0.1$ such that for infinitely many n we have

$$\mathrm{P}\left[\left|\left\langle\mathbf{1}\{\boldsymbol{\sigma}_{x_1} = \omega_1\}|\boldsymbol{\sigma}_{x_2} = \omega_2\right\rangle_{\mu_{\boldsymbol{G}_n}} - \left\langle\mathbf{1}\{\boldsymbol{\sigma}_{x_1} = \omega_1\}\right\rangle_{\mu_{\boldsymbol{G}_n}}\right| > 2\varepsilon,\ \left\langle\mathbf{1}\{\boldsymbol{\sigma}_{x_2} = \omega_2\}\right\rangle_{\mu_{\boldsymbol{G}_n}} > 2\varepsilon\right] > 2\varepsilon. \tag{3.8}$$

Thus, let ℓ be a large enough integer and let $\mathscr{E}$ be the event that the distance between x_1, x_2 in $\boldsymbol{G}_n$ is greater than 2ℓ. Because our factor graph $\boldsymbol{G}_n$ is sparse and random, we have $\mathrm{P}[\mathscr{E}] = 1 - o(1)$ as $n \to \infty$. Therefore, (3.8) implies

$$\mathrm{P}\left[\left|\left\langle\mathbf{1}\{\boldsymbol{\sigma}_{x_1} = \omega_1\}|\boldsymbol{\sigma}_{x_2} = \omega_2\right\rangle_{\mu_{\boldsymbol{G}_n}} - \left\langle\mathbf{1}\{\boldsymbol{\sigma}_{x_1} = \omega_1\}\right\rangle_{\mu_{\boldsymbol{G}_n}}\right| > \varepsilon,\ \left\langle\mathbf{1}\{\boldsymbol{\sigma}_{x_2} = \omega_2\}\right\rangle > \varepsilon,\ \mathscr{E}\right] > \varepsilon. \tag{3.9}$$

To complete the proof, let $\mathscr{S}$ be the set of all $\sigma \in \Omega^{V_n}$ such that $\sigma_{x_2} = \omega_2$. If the event $\mathscr{E}$ occurs, then given $\nabla_\ell(\boldsymbol{G}_n, x_1)$ the value assigned to x_2 is fixed. Therefore, (3.9) implies

$$\mathrm{E}\left\langle\left\|\mu_{\downarrow x_1} - \mu_{\downarrow x_1}[\cdot|\nabla_\ell(\boldsymbol{G}_n, x_1)]\right\|_1\right\rangle_{\mu_{\boldsymbol{G}_n}} \geq \mathrm{E}\left[\mathbf{1}\{\mathscr{E}\}\left\langle\left\|\mu_{\downarrow x_1} - \mu_{\downarrow x_1}[\cdot|\nabla_\ell(\boldsymbol{G}_n, x_1)]\right\|_1|\mathscr{S}\right\rangle_{\mu_{\boldsymbol{G}_n}}\langle\mathbf{1}\{\boldsymbol{\sigma}\in\mathscr{S}\}\rangle_{\mu_{\boldsymbol{G}_n}}\right] \geq \varepsilon^3,$$

in contradiction to (3.6). □

Assumption (3.6), and hence (3.5), is a weaker than the assumption of *Gibbs uniqueness*, another spatial mixing property. Under this stronger assumption Dembo, Montanari, and Sun [18] used an interpolation scheme to compute the free energy in a wide variety of factor models on graphs converging locally to random trees. Further related work on Gibbs uniqueness and/or the interpolation method includes [10, 14, 20].

Although non-reconstruction is sufficient for (3.5) to hold, it is not a necessary condition. Yet there are quite a few examples of random factor graph models where the condition is expected (or known) to hold but where the free energy has not been computed rigorously (e.g., the random k-SAT model [26] or planted models [25]). We expect that these could be tackled via Proposition 3.4 and the other results in this section.

3.4. **Belief Propagation.** To establish a connection between the Gibbs measure $\mu_{\boldsymbol{G}_n}$ and the limit ϑ of the local structure of the factor graph we are going to use the Belief Propagation scheme, which plays a key role in the physicists' "cavity method" [31, Chapter 14]. In fact, due to the Poisson structure of the tree distribution ϑ, Belief Propagation takes a relatively simple form in our setting.

If we look at the random graph $\boldsymbol{G}_n$, then by Proposition 3.3 ϑ gives the fraction of variable nodes x_i such that $\partial^\ell_{\boldsymbol{G}_n} T \cong \partial^\ell T$ for every tree T. Suppose that for each tree T we record the empirical distribution of the marginals $\mu_{\boldsymbol{G}_n\downarrow x_i}$ of such variable nodes. Since each marginal is a distribution on $\mathscr{P}(\Omega)$, this empirical distribution lies in $\mathscr{P}^2(\Omega) = \mathscr{P}(\mathscr{P}(\Omega))$. Hence, we obtain a map $\mathfrak{T} \to \mathscr{P}^2(\Omega)$. According to the Belief Propagation equations, this map must satisfy a certain "consistency condition". To be specific, for a variable node v of T rooted at r_T let $\partial_{T\downarrow} v$ be the set of all children of v. Moreover, let $T_{\downarrow v}$ be the tree "pending on v", i.e., the connected component of v in the tree obtained from T by removing the neighbor a on the path from v to r_T. Then the marginal distribution of T must be "consistent" with the marginal distributions of the trees $T_{\downarrow v}$ for v at distance exactly two from the root r_T.

To formalize this, we call a measurable map $\nu^\star : \mathfrak{T} \to \mathscr{P}^2(\Omega)$, $T \mapsto \nu^\star_T$ a *ϑ-Belief Propagation fixed point* if the following condition holds for ϑ-almost all trees $T \in \mathfrak{T}$. Independently for each variable y at distance precisely two from r_T choose $\eta_{T,y} \in \mathscr{P}(\Omega)$ from the distribution $\nu^\star_{T_{\downarrow y}}$. Moreover, for each constraint node $a \in \partial_T r_T$ let

$$\hat\eta_{T,a}(\omega) \propto \sum_{\sigma \in \Omega^{\partial_T a}} \mathbf{1}\{\sigma_{r_T} = \omega\}\psi_a(\sigma) \prod_{y \in \partial_{T\downarrow} a} \eta_{T,y}(\sigma_y) \qquad (\omega \in \Omega). \tag{3.10}$$

Then we require that

$$\eta_{r_T}(\omega) \propto \prod_{a \in \partial_T r_T} \hat\eta_{T,a}(\omega) \qquad (\omega \in \Omega) \tag{3.11}$$

has distribution $\nu^\star_T$. The idea behind (3.10)–(3.11) is that the marginal distribution of the spin of the root variable behaves as though the the spins assigned to the roots of the subtrees were independent. For a detailed derivation of the Belief Propagation equations see [31, Chapter 14].

We would like to show that under the assumption (3.5) the marginals of the Gibbs measure $\mu_{\boldsymbol{G}_n}$ "converge to" a Belief Propagation fixed point. To this end, we define for a tree $T \in \mathfrak{T}$, an integer $\ell \geq 0$ and a factor graph $\boldsymbol{G}_n$ the distribution $\nu_{\boldsymbol{G}_n,T,\ell} \in \mathscr{P}(\Omega)^2$ as follows. Let $V(\boldsymbol{G}_n, T, \ell)$ be the set of all $x_i \in V_n$ such that $\partial^\ell_{\boldsymbol{G}_n} x_i \cong \partial^\ell T$. If $V(\boldsymbol{G}_n, T, \ell) \neq \emptyset$ we let

$$\nu_{\boldsymbol{G}_n,T,\ell} = \frac{1}{|V(\boldsymbol{G}_n, T, \ell)|} \sum_{x \in V(\boldsymbol{G}_m,T,\ell)} \delta_{\mu_{\boldsymbol{G}_n\downarrow x}}.$$

Thus, $\nu_{\boldsymbol{G}_n,T,\ell}$ is the empirical distribution of the Gibbs marginals $\mu_{\boldsymbol{G}_n\downarrow x}$ for $x \in V(\boldsymbol{G}_n, T, \ell)$. If $V(\boldsymbol{G}_n, T, \ell) = \emptyset$, we let $\nu_{\boldsymbol{G}_n,T,\ell}$ be the uniform distribution on $\mathscr{P}(\Omega)$, say. Recall that $d_1(\cdot,\cdot)$ denotes the L_1-Wasserstein metric.

Theorem 3.5. *If (3.5) holds, then there exists a ϑ-Belief Propagation fixed point $\nu^\star_{\boldsymbol{T}}$ such that*

$$\lim_{\ell\to\infty} \lim_{n\to\infty} \mathrm{E}_{\boldsymbol{G}_n,\boldsymbol{T}}[d_1(\nu_{\boldsymbol{G}_n,\boldsymbol{T},\ell}, \nu^\star_{\boldsymbol{T}})] = 0. \tag{3.12}$$

Thus, for large enough ℓ, n and for a random tree $\boldsymbol{T}$ the empirical distribution of the marginals of those variables of $\boldsymbol{G}_n$ whose depth-ℓ neighborhood is isomorphic to $\partial^\ell \boldsymbol{T}$ is close to $\nu^\star_{\boldsymbol{T}}$. In particular, in the limit $n \to \infty$ the Gibbs measures $\mu_{\boldsymbol{G}_n}$ on the random factor graph induce a Belief Propagation fixed point on the limiting random tree $\boldsymbol{T}$.

Panchenko [33] proved a related fixed point property under different assumptions. While he does not require the convergence in probability assumption (3.5), he imposes certain additional conditions on the random factor graph model (which may not be strictly necessary). Moreover, Panchenko's result just provides a "single-level" distributional fixed point equation, rather than a full "unwrapping" into a ϑ-Belief Propagation fixed point. The proof of Theorem 3.5 is by extension of the argument from [33, Section 2]; in fact, the "unwrapping" requires a fair amount of work.

The proof of Theorem 3.5 yields the following "geometric" fact about the interplay of the marginals on the random factor graph. For a factor graph G and a variable node x let $G - x$ be the factor graph obtained by removing x and its adjacent constraint nodes.

Corollary 3.6. *If (3.5) holds then*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{x \in V_n} \sum_{\omega \in \Omega} \mathrm{E}\left|\mu_{\boldsymbol{G}_n\downarrow x}(\omega) - \frac{\prod_{a \in \partial_{\boldsymbol{G}_n} x} \sum_{\boldsymbol{s}:\partial_{\boldsymbol{G}_n} a \to \Omega} \mathbf{1}\{s_x = \omega\}\psi_a(\boldsymbol{s}) \prod_{y \in \partial_{\boldsymbol{G}_n - x} a} \mu_{\boldsymbol{G}_n - x\downarrow y}(\boldsymbol{s}_y)}{\sum_{\omega' \in \Omega} \prod_{a \in \partial_{\boldsymbol{G}_n} x} \sum_{\boldsymbol{s}:\partial_{\boldsymbol{G}_n} a \to \Omega} \mathbf{1}\{s_x = \omega'\}\psi_a(\boldsymbol{s}) \prod_{y \in \partial_{\boldsymbol{G}_n - x} a} \mu_{\boldsymbol{G}_n - x\downarrow y}(\boldsymbol{s}_y)}\right| = 0,$$

where the expectation is over the choice of the random factor graph $\boldsymbol{G}_n$.

3.5. **The Bethe free energy.** According to the "cavity method", we can extract $\frac{1}{n}\mathrm{E}[\ln Z_{\boldsymbol{G}_n}]$ from the Belief Propagation fixed point from Theorem 3.5 via a formula called the Bethe free energy [31, Section 14.2.4]. Suppose that $\nu^\star$ is a Belief Propagation fixed point. For a tree $T \in \mathfrak{T}$ with root r_T and y at distance two from r_T let $\eta_{T,y} \in \mathscr{P}(\Omega)$ be independently chosen from $\nu^\star_{T_{\downarrow y}}$. Moreover, for $a \in \partial_T r_T$ define $\hat\eta_{T,a} \in \mathscr{P}(\Omega)$ as in (3.10). Further, let

$$\tilde\eta_{T,a}(\sigma) \propto \prod_{b \in \partial_T r_T \setminus \{a\}} \hat\nu_{T,b}(\sigma), \qquad (\sigma \in \Omega), \tag{3.13}$$

$$\varphi_T = \ln \sum_{\sigma \in \Omega} \prod_{a \in \partial_T r_T} \hat\eta_{T,a}(\sigma), \qquad \hat\varphi_{T,a} = \ln \sum_{\sigma \in \Omega^{\partial_T a}} \psi_a(\sigma)\tilde\eta_{T,a}(\sigma_{r_T}) \prod_{y \in \partial_{T\downarrow} a} \eta_{T,y}(\sigma_y), \tag{3.14}$$

$$\tilde\varphi_{T,a} = \ln \sum_{\sigma \in \Omega} \tilde\eta_{T,a}(\sigma)\hat\eta_{T,a}(\sigma). \tag{3.15}$$

(The arguments of all the above logarithms are strictly positive. Indeed, because the functions $\psi \in \Psi$ take strictly positive values, (3.10) ensures that $\hat\eta_{T,a}(\omega) > 0$ for all $\omega \in \Omega$. Hence, φ is well-defined. So are $\hat\varphi_a$, $\tilde\varphi_a$, by the same token.) Taking the expectation over $\boldsymbol{T}$ and independent $\eta_y \in \mathscr{P}(\Omega)$ for all y at distance two from r_T, we define the *Bethe free energy* as

$$\mathscr{B}_\vartheta(\nu^\star_{\boldsymbol{T}}) = \mathrm{E}\left[\varphi_{\boldsymbol{T}} + \sum_{a \in \partial_{\boldsymbol{T}} r_{\boldsymbol{T}}} \left(\frac{\hat\varphi_{\boldsymbol{T},a}}{k_a} - \tilde\varphi_{\boldsymbol{T},a}\right)\right]. \tag{3.16}$$

Theorem 3.7. *Assume that (3.5) holds and let $\nu^\star_{\boldsymbol{T}}$ be a ϑ-Belief Propagation fixed point such that (3.12) holds. Then*

$$\lim_{n\to\infty} \frac{1}{n}\mathrm{E}[\ln Z_{\boldsymbol{G}_n}] = \mathscr{B}_\vartheta(\nu^\star_{\boldsymbol{T}}).$$

Hence, Theorems 3.5 and 3.7 show that under the assumption (3.5) the free energy of the random factor graph model comes out in terms of a "geometric" measure, i.e., a Belief Propagation fixed point on the limiting tree that captures the local structure of the factor graph.

The proof of Theorem 3.7 is by adapting ideas from Panchenko [33, Section 2] to our situation. In particular, we combine the convergence of the measure $\mu_{\boldsymbol{G}_n}$ with a technique from Aizenman, Simms, Starr [4]. The difference between Theorem 3.7 and [33] is that the latter requires certain conditions on the weight functions $\psi \in \Psi$ (to facilitate an interpolation argument) but does not require (3.5). (However, it is stated without proof in [33] that the free energy can be derived along the lines of that paper under the assumption (3.5).)

3.6. **Proof of Theorem 3.5.** We begin by constructing a family of $\mathscr{P}(\Omega)$-valued random variables $(X_\ell)_\ell$ as follows. Let $\boldsymbol{\tau} = (\boldsymbol{\tau}_i)_{i\geq 1}$ be an i.i.d. family of $\mathscr{P}(\Omega)$-valued random variables with distribution $w_{\boldsymbol{x}}$ for a uniformly random $\boldsymbol{x} \in [0,1)$. Then $X_\ell = X_\ell(\boldsymbol{T},\boldsymbol{\tau})$ is defined as follows. We simply set $X_0(\boldsymbol{T},\boldsymbol{\tau}) = \boldsymbol{\tau}_1$. Further, to define $X_\ell(\boldsymbol{T},\boldsymbol{\tau})$ for $\ell \geq 1$ let $V_\ell = \{x_{\ell,1},\ldots,x_{\ell,L}\}$ be the set of variable nodes at distance precisely 2ℓ from the root $r_{\boldsymbol{T}}$, ordered in some arbitrary but deterministic way. Then we let $X_\ell(\boldsymbol{T},\boldsymbol{\tau})$ be the distribution of the spin $\boldsymbol{\sigma}_{r_{\boldsymbol{T}}}$ under the Gibbs measure of the tree $\partial^\ell \boldsymbol{T}$ with a boundary condition chosen independently from $\boldsymbol{\tau}_1,\ldots,\boldsymbol{\tau}_L$. In symbols,

$$X_\ell(\boldsymbol{T},\boldsymbol{\tau}) = \sum_{\sigma \in \Omega^L}\left[\delta_{\langle \boldsymbol{\sigma}_{r_{\boldsymbol{T}}}(\cdot)|\boldsymbol{\sigma}(x_{\ell,1})=\sigma_1,\ldots,\boldsymbol{\sigma}(x_{\ell,L})=\sigma_L\rangle_{\mu_{\partial^\ell \boldsymbol{T}}}} \prod_{i\leq L}\boldsymbol{\tau}_i(\sigma_i)\right] \in \mathscr{P}(\Omega).$$

Now we give an alternative construction. Again set $X_0 = \boldsymbol{\tau}_1$. Now for $\ell \geq 1$, let $X_\ell(\boldsymbol{T},\boldsymbol{\tau})$ be the distribution of the spin at the root given that its neighbors have distribution $X_{\ell-1}(\boldsymbol{T}_1,\boldsymbol{\tau}^1), X_{\ell-1}(\boldsymbol{T}_2,\boldsymbol{\tau}^2),\ldots$ where $\boldsymbol{T}_1$ is the tree appending the ith neighbor of the root of $\boldsymbol{T}$, and $\boldsymbol{\tau}^1,\boldsymbol{\tau}^2,\ldots$ are independent copies of $\boldsymbol{\tau}$. Observe that the two constructions do in fact give the same distribution.

The main step of the proof is to compare X_ℓ with the empirical distribution on the random factor graph. Recall that $d_1(\cdot,\cdot)$ denotes the Wasserstein metric.

Lemma 3.8. *For every $T \in \mathfrak{T}$ the following is true. Let $Y_\ell(T,\boldsymbol{G}_n)$ be the empirical distribution of the marginals of the variables x such that $\partial^\ell_{\boldsymbol{G}_n} x \cong \partial^\ell T$. Then*

$$\lim_{n\to\infty} \mathrm{E}[d_1(Y_\ell(T,\boldsymbol{G}_n), X_\ell(T,\boldsymbol{\tau}))] = 0. \tag{3.17}$$

Proof. The empirical distribution of marginals converges in probability to $\int_0^1 \delta_{w_x}\,dx$ (by assumption), and the distribution of the local neighborhood of x converges to ϑ (Proposition 3.3), but a priori we do not know how the two distributions are coupled. To get a handle on this we proceed by induction on the sub-trees of T. Observe that

there is nothing to show if $\ell = 0$ and T consists of just the root. Now by induction, let us assume that (3.17) holds for any depth-$(\ell-1)$ sub-tree of T.

We need to show that with high probability of the choice of the random factor graph $\boldsymbol{G}_n$, the empirical distribution $Y_\ell(T, \boldsymbol{G}_n)$ converges to that of $X_\ell(T, \boldsymbol{\tau})$. To accomplish this, we will approximate the moments of $Y_\ell(T, \boldsymbol{G}_n)$ and show that these converge to those of $X_\ell(T, \boldsymbol{\tau})$.

Consider the following experiment to approximate the first moment of $Y_\ell(T, \boldsymbol{G}_n)$. Fix T, ℓ, and $\varepsilon > 0$. We choose $L = L(\varepsilon)$ large enough (as we will see below), and sample a random factor graph $\boldsymbol{G}'_{n-L}$ with a slightly sparser constraint density than $\boldsymbol{G}_{n-L}$: instead of adding a Poisson number of constraints ψ with mean $(n-L)\rho_\psi$, we add a Poisson number with mean $n\rho_\psi \cdot ((n-L)/n)^{k_\psi}$. In particular, $\boldsymbol{G}'_{n-L}$ has exactly the distribution of the sub-graph of $\boldsymbol{G}_n$ induced by the first $n-L$ variable nodes. Note that the difference in means is constant, much smaller than $\Theta(\sqrt{n})$, the standard deviation of the number of each constraint type in $\boldsymbol{G}_{n-L}$. The distribution of $\boldsymbol{G}'_{n-L}$ therefore has total variation distance $o(1)$ to $\boldsymbol{G}_{n-L}$, and in particular, the assumption (3.5) holds for $\boldsymbol{G}'_{n-L}$.

We next add to $\boldsymbol{G}'_{n-L}$ L variable nodes $x_{n-L+1}, \ldots, x_n$, along with a Poisson number of each type of constraint node of mean $nk_\psi\rho_\psi(1 - ((n-L)/n)^{k_\psi})$, attached to uniformly random variable nodes from $\boldsymbol{G}_{n-L}$ and the new variable nodes, *conditioned* on the event that at least one of the attached variables nodes of each new constraint node is one of the newly added variable nodes. The resulting factor graph has exactly the distribution of $\boldsymbol{G}_n$.

Now recalling that the tree T is fixed, we condition on the event that the constraint nodes joined to each of the new variable nodes $x_{n-L+1}, \ldots x_n$ are of the number and type joined to the root of T, and their attached variable nodes, call them $x_{11}, \ldots x_{1k}, \ldots x_{L1}, \ldots x_{Lk}$, have depth $\ell-1$ neighborhoods matching those of the attached subtrees $T_1, \ldots T_k$ of T. This event has probability bounded away from 0 as $n \to \infty$ in $\boldsymbol{G}_{n+L}$, and so the resulting conditioned graph, $\boldsymbol{G}'_n$, with the last L variables nodes selected, has the distribution of $\boldsymbol{G}_n$ with L variables nodes selected uniformly at random from all variable nodes whose depth-ℓ neighborhood matches that of T. Let the set $\{x_{i,a,y}\}$, $i = n-L+1, \ldots n$, $a \in \partial_T x_i$, $y \in \partial_{T\downarrow} a$, denote the randomly chosen variable nodes from $\boldsymbol{G}_n$ which are attached to the new constraint nodes in the given positions. Note that whp $\{x_{i,a,y}\}$ will only contain variable nodes from $\boldsymbol{G}'_{n-L}$, as whp no constraint is attached to more than one variable node from a fixed constant-sized set.

Now the graph $\boldsymbol{G}'_n$ induces a marginal distribution on the spin at each of the variable nodes $x_{n-L+1}, \ldots x_n$. Fix $\omega \in \Omega$, and call the ω values of these marginals $q_{n-L+1}(\omega), \ldots, q_n(\omega)$. Let $\overline{q}(\omega) = \frac{1}{L}\sum_{i=n-L+1}^{n} q_i(\omega)$. By choosing $L(\varepsilon)$ large enough, we have that whp over the choice of $\boldsymbol{G}'_{n-L}$ and probability at least $1-\varepsilon$ over the choice of $\{x_{i,a,y}\}$, $\overline{q}(\omega)$ is within ε of the mean of $Y_\ell(T, \boldsymbol{G}_n)(\omega)$.

Similarly, we can approximate the higher moments of $Y_\ell(T, \boldsymbol{G}_n)$, that is, for a vector $\omega_1, \ldots, \omega_r \in \Omega^r$ and powers $i_1, \ldots i_r$, the mean of $\prod_{j=1}^{r} Y_\ell(T, \boldsymbol{G}_n)(\omega_j)^{i_j}$. In this case we again add $L = L(\varepsilon)$ new variable nodes to $\boldsymbol{G}'_{n-L}$ with the appropriate constraint nodes and attached variable nodes and condition that the local neighborhoods of $x_{n-L+1} \ldots x_n$ match T to form $\boldsymbol{G}'_n$. Let $q_i(\omega_j)$ denote the marginal probability of ω_j at variable node x_i, $i = n-L+1, \ldots, n$. Then let $\overline{q} = \frac{1}{L}\sum_{i=n-L+1}^{n}\prod_{j=1}^{r} q_{n+i}(\omega_j)^{i_j}$. Again by choosing $L = L(\varepsilon)$ large enough we can guarantee that whp over the choice of $\boldsymbol{G}'_{n-L}$ and probability at least $1-\varepsilon$ over the choice of $\{x_{i,a,y}\}$, $\overline{q}$ is within ε of the corresponding higher moment of $Y_\ell(T, \boldsymbol{G}_n)$.

What remains to show is that these moment calculations converge to those of $X_\ell(T, \boldsymbol{\tau})$. For $\omega_1, \ldots, \omega_L \in \Omega$ define

$$\boldsymbol{\mu}(\omega_1, \ldots, \omega_L) = \mu_{\boldsymbol{G}'_n}(\{\sigma(x_{n-L+1}) = \omega_1, \ldots, \sigma(x_n) = \omega_L\});$$

then $\boldsymbol{\mu}$ is a random variable, dependent on $\boldsymbol{G}'_{n-L}$ and the family $\{x_{i,a,y}\}$. Hence, if we condition on $\boldsymbol{G}'_{n-L}$, then $\boldsymbol{\mu}(\omega_1, \ldots, \omega_L)$ remains random, namely dependent on $\{x_{i,a,y}\}$. As we saw in the previous paragraph, it suffices to show, for all $\omega_1, \ldots, \omega_L \in \Omega$, that $\boldsymbol{\mu}(\omega_1, \ldots, \omega_L)$ converges in distribution, with high probability over the choice of $\boldsymbol{G}'_{n-L}$, to $\boldsymbol{X}(\omega_1, \ldots, \omega_L) = \prod_{i=1}^{k} \tilde{X}_\ell^{(i)}(T, \boldsymbol{\tau})(\omega_i)$, where the $\tilde{X}_\ell^{(i)}(T, \boldsymbol{\tau})$'s are independent samples from $X_\ell(T, \boldsymbol{\tau})$.

Let Z_{n-L} be the partition function of $\boldsymbol{G}'_{n-L}$, and Z_n the partition function of $\boldsymbol{G}'_n$ in the experiment above. Let $\{x_{i,a,y}\}$, $i = n+1, \ldots n_L$, $a \in \partial_T x_i$, $y \in \partial_{T\downarrow} a$, denote the randomly chosen variable nodes from $\boldsymbol{G}_n$ which are attached to the new constraint nodes in the given positions. Condition on $\boldsymbol{G}'_{n-L}$ and the choice of $\{x_{i,a,y}\}$ and write

$$Z_{n-L} = \sum_{\boldsymbol{s} \in \Omega^{\{x_{i,a,y}\}}} Z_{\boldsymbol{s}}$$

where the vector $\boldsymbol{s}$ represents one set of possible values taken by the variable nodes, and $Z_{\boldsymbol{s}}$ is the partition function of $\boldsymbol{G}_n$ restricted to the set of assignments for which we have $\sigma(x_{i,a,y}) = s_{a,i,y}\ \forall i, a, y$. We can write

$$Z_{\boldsymbol{s}} = Z_{n-L} \cdot \mu_{\boldsymbol{G}_n}(\sigma(\{x_{i,a,y}\}) = \boldsymbol{s}).$$

Given the family $\{x_{i,a,y}\}$, we have

$$\boldsymbol{\mu}(\omega_1,\dots,\omega_L)=\frac{\sum_{\boldsymbol{s}\in\Omega^{\{x_{i,a,y}\}}}\mu_{\boldsymbol{G}'_{n-L}}(\sigma(\{x_{i,a,y}\})=\boldsymbol{s})\prod_{a\in\partial_T x_{n-L+1},\dots\partial_T x_n}\psi_a(\omega_i,\boldsymbol{s}_a)}{\sum_{\omega'_1,\dots\omega'_L\in\Omega^L}\sum_{\boldsymbol{s}\in\Omega^{\{x_{i,a,y}\}}}\mu_{\boldsymbol{G}'_{n-L}}(\sigma(\{x_{i,a,y}\})=\boldsymbol{s})\prod_{a\in\partial_T x_{n-L+1},\dots\partial_T x_n}\psi_a(\omega'_i,\boldsymbol{s}_a)}\,, \tag{3.18}$$

where the quantities on both sides are deterministic numbers, as we have conditioned on $\boldsymbol{G}'_{n-L}$ and the selection of $\{x_{i,a,y}\}$. Now viewing the choice of $\{x_{i,a,y}\}$ as random, we use the asymptotic factorization property (our assumption (3.5) and Corollary 2.8, recalling that this holds whp for $\boldsymbol{G}'_{n-L}$ just as it does for $\boldsymbol{G}_{n-L}$), and get that $\boldsymbol{\mu}(\omega_1,\dots,\omega_L)$ converges in distribution (whp over the choice of $\boldsymbol{G}'_{n-L}$) to

$$\frac{\sum_{\boldsymbol{s}\in\Omega^{\{x_{i,a,y}\}}}\prod_{i=1}^{L}\prod_{a\in\partial_T r_T}\psi_a(\omega_i,\boldsymbol{s}_a)\prod_{y\in\partial_{T\downarrow}a}\tilde{X}_{i,a,y}(\boldsymbol{s}_{i,a,y})}{\sum_{\omega'_1,\dots\omega'_L\in\Omega^L}\sum_{\boldsymbol{s}\in\Omega^{\{x_{i,a,y}\}}}\prod_{i=1}^{L}\prod_{a\in\partial_T r_T}\psi_a(\omega'_i,\boldsymbol{s}_a)\prod_{y\in\partial_{T\downarrow}a}\tilde{X}_{i,a,y}(\boldsymbol{s}_{i,a,y})}\,, \tag{3.19}$$

where the $\tilde{X}_{i,a,y}$'s are independent samples from the respective distributions $X_{\ell-1}(T_{a,y},\boldsymbol{\tau})$. We rearrange (3.19) to give the following convergence in distribution, whp over the choice of $\boldsymbol{G}'_n$:

$$\boldsymbol{\mu}(\omega_1,\dots,\omega_L)\Rightarrow\prod_{i=1}^{L}\frac{\prod_{a\in\partial_T r_T}\sum_{\boldsymbol{s}\in\Omega^{k_a}}\mathbf{1}_{s_j=\omega_i}\cdot\psi_a(\boldsymbol{s})\prod_{y\in\partial_{T\downarrow}a}\tilde{X}_{i,a,y}(\boldsymbol{s}_y)}{\sum_{\omega'\in\Omega}\prod_{a\in\partial_T r_T}\sum_{\boldsymbol{s}\in\Omega^{k_a}}\mathbf{1}_{s_j=\omega'}\cdot\psi_a(\boldsymbol{s})\prod_{y\in\partial_{T\downarrow}a}\tilde{X}_{i,a,y}(\boldsymbol{s}_y)}\,. \tag{3.20}$$

Finally, performing the calculation (3.19) in reverse on L independent copies of the tree T, we see that the r.h.s. of (3.20) is distributed as $\boldsymbol{X}(\omega_1,\dots,\omega_L)$. In particular this holds for every fixed L and choice of $\omega_1,\dots\omega_L$, and so this proves convergence of the moments. □

To obtain the desired Belief Propagation fixed point we let $\nu_{T,\ell}\in\mathscr{P}^2(\Omega)$ be the distribution of $X_\ell(T,\boldsymbol{\tau})$. Moreover, let $\mathscr{F}_\ell$ be the σ-algebra on $\mathfrak{T}$ generated by events $\{\partial^\ell\boldsymbol{T}\cong\partial^\ell T\}$ for $T\in\mathfrak{T}$.

Corollary 3.9. *For any $\ell\geq 0$ we have* $\mathrm{E}[\nu_{\boldsymbol{T},\ell+1}|\mathscr{F}_\ell]=\nu_{\boldsymbol{T},\ell}$.

Proof. As a first step we are going to show that

$$\mathrm{E}\left[\nu_{\boldsymbol{T},1}|\mathscr{F}_0\right]=\mathrm{E}\left[\nu_{\boldsymbol{T},1}\right]=\nu_{\boldsymbol{T},0}. \tag{3.21}$$

Observe that by the definition of X_0 the right-hand side is in deterministic. To prove (3.21) we apply Lemma 3.8 to $\ell=1$. Our assumption (3.5) implies that the average empirical distribution $\mathrm{E}[Y_1(\boldsymbol{T},\boldsymbol{G}_n)]$ converges to the distribution of $w_{\boldsymbol{x}}$ for a uniform $\boldsymbol{x}\in[0,1)$ as $n\to\infty$. The latter is precisely the law of $\nu_{\boldsymbol{T},0}$. Moreover, $\mathrm{E}[Y_1(\boldsymbol{T},\boldsymbol{G}_n)]$ converges to $\mathrm{E}[\nu_{\boldsymbol{T},1}]$ by Lemma 3.8. Therefore, $\mathrm{E}[\nu_{\boldsymbol{T},1}]=\nu_{\boldsymbol{T},0}$.

For general values of ℓ we proceed by induction. If we condition on the first 2ℓ levels of the random tree $\boldsymbol{T}$, then the trees pending on the variable nodes at distance precisely 2ℓ from the root are independent copies of $\boldsymbol{T}$ itself. Therefore, the levels $2\ell+1$ and $2\ell+2$ are distributed as $\partial^1\boldsymbol{T}$. Hence, let V_ℓ be the set of variable nodes at distance precisely 2ℓ from the root. Then (3.21) implies that the distribution of each $x\in V_\ell$ under a random boundary condition for $V_{\ell+1}$ as in the construction of $X_{\ell+1}$ is identical to the distribution of $w_{\boldsymbol{x}}$. Further, all $x\in V_\ell$ are independent. Consequently, $\mathrm{E}[\nu_{\boldsymbol{T},\ell+1}|\mathscr{F}_\ell]=\nu_{\boldsymbol{T},\ell}$. □

Proof of Theorem 3.5. Corollary 3.9 implies that for any continuous $f:\mathscr{P}(\Omega)\to\mathbb{R}$ the sequence $(\int f\,d\nu_{\boldsymbol{T},\ell})_\ell$ is a martingale w.r.t. $(\mathscr{F}_\ell)_\ell$. Because it is bounded, the martingale converges almost surely and in L_1. Furthermore, because the space of $C(\mathscr{P}(\Omega))$ of continuous functions on $\mathscr{P}(\Omega)$ has a countable dense set with respect to uniform convergence (e.g., the polynomials with rational coefficients by Weierstrass), we find that for almost all $\boldsymbol{T}$ the sequences $(\int f\,d\nu_{\boldsymbol{T},\ell})_\ell$ converge for all $f\in C(\mathscr{P}(\Omega))$. Given that this event occurs, the map $\nu^\star_{\boldsymbol{T}}:f\mapsto\lim_{\ell\to\infty}\int f\,d\nu_{\boldsymbol{T},\ell}$ is a continuous linear functional on $C(\mathscr{P}(\Omega))$ that satisfies $\nu^\star_{\boldsymbol{T}}(1)=1$. Hence, by the Riesz representation theorem $\nu^\star_{\boldsymbol{T}}$ is a probability measure on $\mathscr{P}(\Omega)$. Furthermore, our definition (3.3) of the topology of $\mathfrak{T}$ ensures that for each ℓ the function $T\mapsto\nu_{T,\ell}$ is continuous. Therefore, being a pointwise limit of the continuous functions, the function $T\mapsto\nu^\star_{\boldsymbol{T}}$ is measurable.

To establish the fixed point property, we recall that Lemma 3.8 and (3.20) imply the following. For a tree T and a constraint $a\in\partial_T r_T$ and let $(\tilde{X}_{a,y})_y$ be a family of independent copies of $X_{\ell-1}(T_{\downarrow y})$ for $y\in\partial_{T\downarrow}a$. Then

$$X_\ell(T,\boldsymbol{\tau})\stackrel{\mathrm{d}}{=}\frac{\prod_{a\in\partial_T r_T}\sum_{\boldsymbol{s}\in\Omega^{k_a}}\mathbf{1}_{s_j=\omega_i}\cdot\psi_a(\boldsymbol{s})\prod_{y\in\partial_{T\downarrow}a}\tilde{X}_{a,y}(\boldsymbol{s}_y)}{\sum_{\omega'\in\Omega}\prod_{a\in\partial_T r_T}\sum_{\boldsymbol{s}\in\Omega^{k_a}}\mathbf{1}_{s_j=\omega'}\cdot\psi_a(\boldsymbol{s})\prod_{y\in\partial_{T\downarrow}a}\tilde{X}_{a,y}(\boldsymbol{s}_y)}$$

Hence, taking the limit $\ell\to\infty$, we conclude that $\nu^\star_{\boldsymbol{T}}$ is a Belief Propagation fixed point. □

Proof of Corollary 3.6. Corollary 3.6 follows from the proof above, by taking $L = 1$ in the construction, then from (3.18) using Corollary 2.8 to show that the joint probability $\mu_{\boldsymbol{G}'_{n-1}}(\sigma(\{x_{a,y}\}) = \boldsymbol{s})$ asymptotically factorizes. Here we do not need to use the fact that the distribution of each marginal converges to the same thing, simply the factorization property. □

3.7. **Proof of Theorem 3.7.** Throughout this section we keep the notation and the assumptions of Theorem 3.7. Moreover, we let $\boldsymbol{x} = (\boldsymbol{x}_i)_{i\geq 1}$ be a sequence of uniform random variables on $[0,1)$ that are mutually independent and independent of everything else. We begin with two claims regarding the average marginals of the fixed point and the Bethe free energy.

Claim 3.10. *Let $\boldsymbol{x} \in [0,1)$ be uniform. Then $\mathrm{E}[\nu^\star_{\boldsymbol{T}}] = \mathscr{L}(w_{\boldsymbol{x}})$.*

Proof. Assumption (3.5) implies that the empirical distribution of the marginals of $\boldsymbol{G}_n$ converges in distribution to $\mathscr{L}(w_{\boldsymbol{x}})$. Hence, Proposition 3.3 and the assumption that $\nu^\star$ satisfies (3.12) imply the assertion. □

Claim 3.11. *Let $(\hat{\boldsymbol{\eta}}_{\psi,j,i})_{\psi\in\Psi, j\in[k_\psi], i\geq 1}$ be a family of independent $\mathscr{P}(\Omega)$-valued random variables with distribution*

$$\hat{\boldsymbol{\eta}}_{\psi,j,i}(\sigma_j) \propto \sum_{(\sigma_h)_h \in \Omega^{[k_\psi]\setminus\{j\}}} \psi(\sigma) \prod_{h\in[k_\psi]\setminus\{j\}} w_{\boldsymbol{x}_h}(\sigma_h) \tag{3.22}$$

Further, let $(d_{\psi,j})_{\psi\in\Psi, j\in[k_\psi]}$ be a family of independent random variables such that $d_{\psi,j}$ has distribution $\mathrm{Po}(\rho_\psi/j)$. Moreover, let $(\boldsymbol{\eta}_h)_{h\geq 1}$ be a family of independent $\mathscr{P}(\Omega)$-valued random variables with distribution

$$\boldsymbol{\eta}_h(\sigma) \propto \prod_{\psi\in\Psi, j\in[k_\psi], i\in[d_{\psi,j}]} \hat{\boldsymbol{\eta}}_{\psi,j,i}(\sigma).$$

Further, let

$$\phi = \ln \sum_{\sigma\in\Omega} \prod_{\psi,j,i} \hat{\boldsymbol{\eta}}_{\psi,j,i}(\sigma),$$
$$\hat{\phi}_\psi = \ln \sum_{\sigma\in\Omega^{k_\psi}} \psi(\sigma) \prod_{h\in[k_\psi]} \boldsymbol{\eta}_h(\sigma_h), \qquad \tilde{\phi}_{\psi,j} = \ln \sum_{\sigma\in\Omega} \boldsymbol{\eta}_1(\sigma)\hat{\boldsymbol{\eta}}_{\psi,j,1}(\sigma) \qquad (\psi\in\Psi,\ j\in[k_\psi]).$$

Then

$$\mathscr{B}_\vartheta(\nu^\star) = \mathrm{E}[\phi] + \sum_{\psi\in\Psi}\sum_{j\in[k_\psi]} \rho_\psi\left[k_\psi^{-1}\mathrm{E}[\hat{\phi}_\psi] - \mathrm{E}[\tilde{\phi}_{\psi,j}]\right]. \tag{3.23}$$

Proof. To show that (3.23) and (3.16) coincide, we recall the random variables from (3.13)–(3.15). Comparing (3.16) and (3.23), we see that it suffices to show

$$\mathrm{E}[\phi] = \mathrm{E}[\varphi_{\boldsymbol{T}}], \qquad \sum_{\psi\in\Psi, j\in[k_\psi]} \rho_\psi \mathrm{E}[\hat{\phi}_{\psi,j}] = \mathrm{E}\left[\sum_{a\in\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}} \hat{\varphi}_{\boldsymbol{T},a}\right], \qquad \sum_{\psi\in\Psi, j\in[k_\psi]} \rho_\psi \mathrm{E}[\tilde{\phi}_{\psi,j}] = \mathrm{E}\left[\sum_{a\in\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}} \tilde{\varphi}_{\boldsymbol{T},a}\right]. \tag{3.24}$$

To prove the first equality, we recall that by construction the root of the random tree $\boldsymbol{T}$ has $\mathrm{Po}(\rho_\psi)$ children a with weight function ψ. For each of them the position $j \in [k_\psi]$ such that $\partial_{\boldsymbol{T}}(a,j) = r_{\boldsymbol{T}}$ is uniform, and they are independent. Therefore, for each j the number of a with $\partial_{\boldsymbol{T}}(a,j) = r_{\boldsymbol{T}}$ has distribution $\mathrm{Po}(\rho_\psi/k_\psi)$, and these random variables are independent. Hence, the distribution of the offspring of the root of $\boldsymbol{T}$ coincides with the joint distribution of the random variables $(d_{\psi,j})_{\psi,j}$. Further, for each $a \in \partial_{\boldsymbol{T}} r_{\boldsymbol{T}}$ and every $y \in \partial_{\boldsymbol{T}} a$ the tree $\boldsymbol{T}_{\downarrow y}$ pending on y is an independent copy of $\boldsymbol{T}$. Therefore, Claim 3.10 implies that the distribution of the random variables $\eta_{\boldsymbol{T},y}$ that go into (3.10) coincides with the distribution of $w_{\boldsymbol{x}}$ for a uniform $\boldsymbol{x} \in [0,1)$. Moreover, the random variables $\eta_{\boldsymbol{T},y}$ are mutually independent. Consequently, comparing (3.22) and (3.10), we see the distribution of $\hat{\eta}_{\boldsymbol{T},a}$ that for a constraint node a with $\psi_a = \psi$ and $\partial_{\boldsymbol{T}}(a,j) = r_{\boldsymbol{T}}$ coincides with the distribution of $\hat{\boldsymbol{\eta}}_{\psi,j,i}(\sigma_j)$ for any $i \geq 1$. Because the $(\hat{\eta}_{\boldsymbol{T},a})_{a\in\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}}$ are mutually independent, we thus see that $\varphi_{\boldsymbol{T}}$ has the same distribution as ϕ. In particular, $\mathrm{E}[\varphi_{\boldsymbol{T}}] = \mathrm{E}[\phi]$.

Further, we derive the middle equation from the Chen-Stein property (1.1). Specifically, remembering that we picked one specific representative of each isomorphism class, we let $a_{\boldsymbol{T}}$ be a random neighbor of the root of $\boldsymbol{T}$. Then

$$\mathrm{E}\left[\sum_{a\in\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}} \hat{\varphi}_{\boldsymbol{T},a}\right] = \mathrm{E}\left[|\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}| \hat{\varphi}_{\boldsymbol{T},a_{\boldsymbol{T}}}\right]. \tag{3.25}$$

Moreover, by a similar argument as in the previous paragraph the random variables $\eta_{\boldsymbol{T},y}$ for y at distance two from $r_{\boldsymbol{T}}$ are independent copies of $w_{\boldsymbol{x}}$. Hence, the $\tilde{\eta}_{\boldsymbol{T},a_{\boldsymbol{T}}}$ factor consists of $|\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}| - 1$ independent factors. By comparison, the terms $\boldsymbol{\eta}_h$ comprise of $|\partial_{\boldsymbol{T}} r_{\boldsymbol{T}}|$ independent factors. Therefore, applying (1.1) to (3.25), we obtain the middle equation of (3.24). The last equation follows form a similar argument. □

To derive the theorem from Claim 3.11 we employ the Aizenman-Simms-Starr scheme [4], i.e., the observation that

$$\frac{1}{n}\mathrm{E}[\ln Z_{\boldsymbol{G}_n}] = \frac{1}{n}\sum_{h=1}^{n} \mathrm{E}\ln \frac{Z_{\boldsymbol{G}_h}}{Z_{\boldsymbol{G}_{h-1}}} \tag{3.26}$$

(with the convention that $\boldsymbol{G}_0$ is the empty graph and $Z_{\boldsymbol{G}_0} = 1$). The assertion follows from (3.26) if we can show that

$$\lim_{n\to\infty} \mathrm{E}\ln \frac{Z_{\boldsymbol{G}_{n+1}}}{Z_{\boldsymbol{G}_n}} = \mathscr{B}_\vartheta(\nu^\star). \tag{3.27}$$

Hence, we need to compare $\boldsymbol{G}_{n+1}, \boldsymbol{G}_n$ for large n. To this end, we couple the two random graphs as follows. Let $\rho'_\psi = (n/(n+1))^{k_\psi}\rho_\psi$ and let $\boldsymbol{G}'$ be a random factor graph on the variable set V_n obtained by including $m'_\psi = \mathrm{Po}(\rho'_\psi n)$ random constraints of type ψ for each $\psi \in \Psi$. Then the distribution of $\boldsymbol{G}'$ coincides with the distribution of $\boldsymbol{G}_{n+1}$ with variable x_{n+1} and all incident constraints removed. Further, obtain $\boldsymbol{G}''$ by adding $m''_\psi = \mathrm{Po}((\rho_\psi - \rho'_\psi)n)$ random constraints to $\boldsymbol{G}'$ independently for each $\psi \in \Psi$. Then $\boldsymbol{G}''$ is distributed as $\boldsymbol{G}_n$. Finally, obtain $\boldsymbol{G}'''$ from $\boldsymbol{G}'$ by adding one variable x_{n+1} and $m'''_\psi = \mathrm{Po}(k_\psi\rho_\psi)$ random constraints of type ψ that are incident to x_{n+1} independently for each $\psi \in \Psi$. Then the distribution of $\boldsymbol{G}'''$ matches that of $\boldsymbol{G}_{n+1}$.

To calculate $\mathrm{E}\ln(Z_{\boldsymbol{G}''}/Z_{\boldsymbol{G}'})$, fix some $\varepsilon > 0$. There is a number $c > 0$ such that $\max\{|\ln\psi(\sigma)| : \psi \in \Psi, \sigma \in \Omega^{k_\psi}\} < c$. Hence, because the tails of the Poisson decay sub-exponentially, there exists $L = L(\varepsilon) > 0$ such that

$$\mathrm{E}\left[\mathbf{1}\{\textstyle\sum_\psi m'''_\psi > L\}|\ln(Z_{\boldsymbol{G}'''}/Z_{\boldsymbol{G}'})|\right] < \varepsilon. \tag{3.28}$$

We are now going to insert the new constraints into $\boldsymbol{G}'$ one by one and track the impact of each insertion. Number the new constraints in some way as $a''_1, \ldots, a''_l$ with $l \le L$, let $\boldsymbol{G}''_0 = \boldsymbol{G}'$ and let $\boldsymbol{G}''_i$ be the obtained by inserting $a''_1, \ldots, a''_i$. Because the total number of constraints that we add is of lower order than the standard variation $\Theta(\sqrt{n})$ of the number of constraints of each type, the total variation distance of $\boldsymbol{G}''_i$ and $\boldsymbol{G}_n$ is $o(1)$ for each $i \le L$. Therefore, if we let $(x_{ij})_{j\in[k_i]}$ be the family of independent, uniformly chosen neighbors of a''_i, then Corollary 2.8 and our assumption (3.5) imply that with high probability

$$\left\|\mu_{\boldsymbol{G}''_{i-1}\downarrow\{x_{ij}:j\in[k_i]\}} - \bigotimes_{j\in[k_i]} \mu_{\boldsymbol{G}''_{i-1}\downarrow x_{ij}}\right\|_1 = o(1) \qquad \text{as } n\to\infty.$$

Together with (3.5) this implies that the vector $(\mu_{\boldsymbol{G}'_{i-1}\downarrow x_{ij}})_{j\in[k_\psi]}$ converges in distribution to $(w_{\boldsymbol{x}_j})_{j\in[k_\psi]}$ for independent uniform $\boldsymbol{x}_j \in [0,1)$. Hence,

$$\mathrm{E}\ln\left(Z_{\boldsymbol{G}''_i}/Z_{\boldsymbol{G}''_{i-1}}\right) = \mathrm{E}\left[\ln \sum_{\sigma\in\Omega^{k_i}} \psi_{a''_i}(\sigma)\prod_{j\in[k_1]} w_{\boldsymbol{x}_j}(\sigma_j)\right] + o(1). \tag{3.29}$$

Summing (3.29) up over $l \le L$, using (3.28) and letting $\varepsilon \to 0$ sufficiently slowly, we obtain

$$\mathrm{E}\ln(Z_{\boldsymbol{G}''}/Z_{\boldsymbol{G}'}) = \sum_{\psi\in\Psi}(\rho_\psi - \rho'_\psi)\mathrm{E}\left[\ln\sum_{\sigma\in\Omega^{k_\psi}}\psi(\sigma)\prod_{j\in[k_\psi]} w_{\boldsymbol{x}_j}(\sigma_j)\right] + o(1). \tag{3.30}$$

Finally, in the notation of Claim 3.11, (3.30) reads

$$\mathrm{E}\ln(Z_{\boldsymbol{G}'''}/Z_{\boldsymbol{G}'}) = \sum_{\psi\in\Psi}(\rho_\psi - \rho'_\psi)\mathrm{E}\left[\hat{\phi}_\psi\right] + o(1). \tag{3.31}$$

To calculate $\mathrm{E}\ln(Z_{\boldsymbol{G}'''}/Z_{\boldsymbol{G}'})$ let $a_1, \ldots, a_l$ be the new constraints attached to x_{n+1}. With probability $1 - O(1/n)$ the new variable x_{n+1} appears precisely once in each a_i. If so, then by the same token as in the previous paragraph the joint distribution $\mu_{\boldsymbol{G}'\downarrow Y}$ of the variables $Y = \bigcup_{i\le l}\partial_{\boldsymbol{G}'''} a_i \setminus \{x_{n+1}\}$ converges to $(w_{\boldsymbol{x}_y})_{y\in Y}$ with independent uniform $\boldsymbol{x}_y \in [0,1)$. Consequently,

$$\mathrm{E}\ln(Z_{\boldsymbol{G}'''}/Z_{\boldsymbol{G}'}) = o(1) + \mathrm{E}\left[\ln\sum_{\sigma\in\Omega^{\{x_{n+1}\}\cup Y}}\prod_{i=1}^{l}\left(\psi_{a_i}((\sigma_y)_{y\in\partial_{\boldsymbol{G}'''}a_i})\prod_{y\in Y\cap\partial_{\boldsymbol{G}'''}a_i} w_{\boldsymbol{x}_y}(\sigma_y)\right)\right]. \tag{3.32}$$

We introduce probability distributions on Ω by letting

$$\hat{\nu}_i(\sigma_{x_{n+1}}) \propto \sum_{(\sigma_y)_{y\in\Omega^{\partial \boldsymbol{G}''' a_i\setminus\{x_{n+1}\}}}} \psi_{a_i}(\sigma) \prod_{y\in\partial_{\boldsymbol{G}''' a_i\setminus\{x_{n+1}\}}} w_{\boldsymbol{x}_y}(\sigma_y), \qquad \nu(\sigma_{x_{n+1}}) \propto \prod_{i=1}^{l}\hat{\nu}_i(\sigma_{x_{n+1}}) \qquad (\sigma_{x_{n+1}}\in\Omega).$$

Then (3.32) becomes

$$\begin{aligned}\mathrm{E}\ln(Z_{\boldsymbol{G}'''}/Z_{\boldsymbol{G}'})+o(1) &= \mathrm{E}\left[\ln\sum_{\sigma_{x_{n+1}}}\prod_{i=1}^{l}\left[\hat{\nu}_i(\sigma_{x_{n+1}})\sum_{\tau_{x_{n+1}}}\hat{\nu}_i(\tau_{x_{n+1}})\right]\right]\\ &= \mathrm{E}\left[\ln\sum_{\sigma_{x_{n+1}}}\prod_{i=1}^{l}\hat{\nu}_i(\sigma_{x_{n+1}})\right]+\sum_{i=1}^{l}\mathrm{E}\left[\ln\sum_{\tau_{x_{n+1}}}\hat{\nu}_i(\tau_{x_{n+1}})\right].\end{aligned} \tag{3.33}$$

Further,

$$\ln\sum_{\tau_{x_{n+1}}}\hat{\nu}_i(\tau_{x_{n+1}}) = \ln\left[\sum_{\tau\in\Omega^{\partial \boldsymbol{G}''' a_i}}\psi_{a_i}(\tau)\nu(\tau_{x_{n+1}})\prod_{y\in\partial_{\boldsymbol{G}''' a_i\setminus\{x_{n+1}\}}}w_{\boldsymbol{x}_y}(\tau_y)\right]-\ln\sum_{\tau_{x_{n+1}}}\nu(\tau_{x_{n+1}})\hat{\nu}_i(\tau_{x_{n+1}}). \tag{3.34}$$

To connect these formulas with Lemma 3.11, we observe that $\sum_{\sigma_{x_{n+1}}}\prod_{i=1}^{l}\hat{\nu}_i(\sigma_{x_{n+1}})$ is distributed as ϕ from Lemma 3.11. Moreover, by the Chen-Stein property we have

$$\mathrm{E}\left[\sum_{i=1}^{l}\ln\sum_{\tau}\psi_{a_i}(\tau)\nu(\tau_{x_{n+1}})\prod_{y\in\partial_{\boldsymbol{G}''' a_i}}w_{\boldsymbol{x}_y}(\tau_y)\right] = \sum_{\psi\in\Psi}k_\psi\rho_\psi\mathrm{E}[\hat{\phi}_\psi],$$

$$\mathrm{E}\left[\sum_{i=1}^{l}\ln\sum_{\tau_{x_{n+1}}}\nu(\tau_{x_{n+1}})\hat{\nu}_i(\tau_{x_{n+1}})\right] = \sum_{\psi\in\Psi}\sum_{j\in[k_\psi]}\rho_\psi\mathrm{E}[\tilde{\phi}_{\psi,j}].$$

Therefore,

$$\mathrm{E}\ln(Z_{\boldsymbol{G}'''}/Z_{\boldsymbol{G}'}) = \mathrm{E}[\phi]+\sum_{\psi\in\Psi}k_\psi\rho_\psi\mathrm{E}[\hat{\phi}_\psi]+\sum_{\psi\in\Psi,j\in[k_\psi]}\rho_\psi\mathrm{E}[\tilde{\phi}_{\psi,j}]+o(1). \tag{3.35}$$

Finally, combining (3.31) and (3.35), we obtain (3.27).

REFERENCES

[1] D. Achlioptas, C. Moore: Random k-SAT: two moments suffice to cross a sharp threshold. SIAM Journal on Computing **36** (2006) 740–762.
[2] D. Achlioptas, A. Naor, and Y. Peres: Rigorous location of phase transitions in hard optimization problems. Nature **435** (2005) 759–764.
[3] D. Achlioptas, A. Naor, Y. Peres: On the maximum satisfiability of random formulas. J. ACM **54** (2007).
[4] M. Aizenman, R. Sims, S. Starr: An extended variational principle for the SK spin-glass model. Phys. Rev. B **68** (2003) 214403.
[5] D. Aldous: Representations for partially exchangeable arrays of random variables. J. Multivariate Anal. **11** (1981) 581–598.
[6] D. Aldous, R. Lyons: Processes on unimodular random networks. Electron. J. Probab. **12** (2007) 1454–1508.
[7] T. Austin: Exchangeable random measures. Ann. Inst. H. Poincaré Probab. Statist. **51** (2015), 842-861.
[8] V. Bapst, A. Coja-Oghlan: Harnessing the Bethe free energy. arXiv:1504.03975 (2015).
[9] V. Bapst, A. Coja-Oghlan: The condensation phase transiontion in the regular k-SAT model. arXiv:1507.03512 (2015).
[10] M. Bayati, D. Gamarnik, P. Tetali: Combinatorial approach to the interpolation method and scaling limits in sparse random graphs. Annals of Probability **41** (2013) 4080–4115.
[11] N. Bhatnagar, A. Sly, P. Tetali: Reconstruction threshold for the hardcore model. Proc. 14th RANDOM (2010) 434–447.
[12] C. Bordenave, P. Caputo: Large deviations of empirical neighborhood distribution in sparse random graphs. Probability Theory and Related Fields **163** (2015) 149–222.
[13] C. Borgs, J. Chayes, L. Lovász, V. Soś, K. Vesztergombi: Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. Advances in Mathematics **219** (2008) 1801–1851.
[14] V. Chandrasekaran, M. Chertkov, D. Gamarnik, D. Shah, J. Shin: Counting independent sets using the Bethe approximation. SIAM J. Discrete Mathematics **25** (2011) 1012–1034.
[15] L. Chen: Poisson approximation for dependent trials. Annals of Probability (1975): 534–545.
[16] C. Efthymiou: Reconstruction/non-reconstruction thresholds for colourings of general Galton-Watson trees. Proc. 19th RANDOM (2015) 756–774.
[17] A. Coja-Oghlan, C. Efthymiou, N. Jaafari: Local convergence of random graph colorings. Proc. 19th RANDOM (2015) 726–737.
[18] A. Dembo, A. Montanari, N. Sun: Factor models on locally tree-like graphs. Annals of Probability **41** (2013) 4162–4213.
[19] P. Diaconis, S. Janson: Graph limits and exchangeable random graphs. Rend. Mat. Appl. **28** (2008) 33–61.
[20] D. Gamarnik, T. Nowicki, G. Swirszcz: Maximum weight independent sets and matchings in sparse random graphs. Exact results using the local weak convergence method. Random Structures & Algorithms **28** (2006) 76–106.
[21] D. Hoover: Relations on probability spaces and arrays of random variables. Preprint, Institute of Advanced Studies, Princeton, 1979.
[22] S. Janson: Graphons, cut norm and distance, couplings and rearrangements. NYJM Monographs, Volume 4, 2013.

[23] S. Janson, T. Łuczak, A. Ruciński: Random Graphs, Wiley 2000.
[24] O. Kallenberg: Probabilistic symmetries and invariance principles. Springer, New York, 2005.
[25] F. Krzakala, L. Zdeborová: Hiding quiet solutions in random constraint satisfaction problems. Phys. Rev. Lett. **102** (2009) 238701.
[26] F. Krzakala, A. Montanari, F. Ricci-Tersenghi, G. Semerjian, L. Zdeborová: Gibbs states and the set of solutions of random constraint satisfaction problems. Proc. National Academy of Sciences **104** (2007) 10318–10323.
[27] D. Kunszenti-Kovács, L. Lovász, B. Szegedy: Multigraph limits, unbounded kernels, and Banach space decorated graphs. arXiv:1406.7846 (2014)
[28] D. Levin, Y. Peres, E. Wilmer: Markov chains and mixing times. AMS 2009.
[29] L. Lovász: Large networks and graph limits. Colloquium Publications **60** (2012), AMS.
[30] L. Lovász, B. Szegedy: Limits of compact decorated graphs. arXiv 1010.5155 (2010).
[31] M. Mézard, A. Montanari: Information, physics and computation. Oxford University Press 2009.
[32] E. Mossel, J. Neeman, A. Sly: Reconstruction and estimation in the planted partition model. Probability Theory and Related Fields (2014) 1–31.
[33] D. Panchenko: Spin glass models from the point of view of spin distributions. Annnals of Probability, **41** (2013) 1315–1361.
[34] D. Panchenko: The Sherrington-Kirkpatrick model. Springer 2013.
[35] T. Richardson, R. Urbanke: Modern coding theory. Cambridge University Press (2008).
[36] E. Szemerédi: Regular partitions of graphs. Colloq. Internat. CNRS **260** (1978) 399–401.
[37] T. Tao: Szemerédi's regularity lemma revisited. Contrib. Discrete Math. **1** (2006) 8–28.
[38] J. Yedidia, W. Freeman, Y. Weiss: Constructing free-energy approximations and generalized Belief Propagation algorithms. IEEE Transactions on Information Theory **51** (2005) 2282–2312.

AMIN COJA-OGHLAN, acoghlan@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.

WILL PERKINS, math@willperkins.org, SCHOOL OF MATHEMATICS, UNIVERSITY OF BIRMINGHAM, EDGBASTON, BIRMINGHAM, UK.

KATHRIN SKUBCH, skubch@math.uni-frankfurt.de, GOETHE UNIVERSITY, MATHEMATICS INSTITUTE, 10 ROBERT MAYER ST, FRANKFURT 60325, GERMANY.